



Identifying Mechanisms of Apoptotic Pore Formation With Programmatic Ensemble Modeling

Citation

Bachman, John. 2016. Identifying Mechanisms of Apoptotic Pore Formation With Programmatic Ensemble Modeling. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:26718762>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Identifying mechanisms of apoptotic pore formation with programmatic ensemble modeling

A DISSERTATION PRESENTED

BY

JOHN ATA BACHMAN

TO

THE COMMITTEE ON HIGHER DEGREES IN SYSTEMS BIOLOGY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN THE SUBJECT OF

SYSTEMS BIOLOGY

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS

DECEMBER 2015

©2015 – JOHN ATA BACHMAN
ALL RIGHTS RESERVED.

Identifying mechanisms of apoptotic pore formation with programmatic ensemble modeling

ABSTRACT

Mitochondrial outer membrane permeabilization is a key step in the apoptotic cell death program, regulating life-death decisions in response to cytotoxic drugs and other forms of cell stress. In this thesis I use mathematical modeling of a reconstituted biochemical system to identify and integrate mechanisms of apoptotic pore formation. A key bottleneck in using mathematical models to characterize mechanisms has been the difficulty of efficiently creating and revising alternative models and evaluating them against data. This problem is addressed through the use of a software framework, PySB, that allows ensembles of models to be transparently described using tools and approaches from computer programming. These alternative hypotheses can then be evaluated against data using methods from Bayesian statistics for discrimination of models with varying numbers of (possibly non-identifiable) parameters. Using this framework, calibration of a set of models to *in vitro* kinetic measurements of the membrane insertion of Bax identifies a conformational intermediate associated with BH₃-only:Bax complex formation and membrane association but not pore formation. Functional measurements of Bax point mutants from tumors show that the pore formation process can be blocked at the transition into or out of this intermediate, preventing mitochondrial permeabilization. In another study, model-based analysis of Bax insertion and permeabilization kinetics across a range of BH₃-only, Bax, and liposome concentrations reveals the context-dependence of the mechanisms regulating pore formation. Bax recruitment is shown to depend on liposome concentration kinetically but not stoichiometrically, whereas cBid recruitment is shown to be limited at high cBid:liposome concentrations. I show that Bax distribution among liposomes is dependent on the presence of pre-existing pores, and that pores grow to include large numbers of Bax monomers but have a minimum size of four subunits. More generally, these studies serve as examples of how ensemble modeling can be used to integrate information about complex mechanisms from disparate sets of experimental observations.

Contents

1	INTRODUCTION	I
1.1	Mathematical models integrate knowledge about mechanisms	3
1.2	Programmatic ensemble modeling	4
1.3	Molecular mechanisms in apoptosis	8
1.4	Systems biochemistry allows quantitative characterization of mechanisms	12
1.5	Synopsis of thesis	13
2	PROGRAMMING BIOLOGICAL MODELS IN PYTHON USING PySB	24
2.1	Introduction	25
2.2	Results	31
2.3	Discussion	48
2.4	Methods	55
2.5	Supplementary information	66
3	CANCER-ASSOCIATED BAX POINT MUTATIONS BLOCK APOPTOTIC PORE FORMATION AT DISTINCT CONFORMATIONS	85
3.1	Introduction	87
3.2	Results	89
3.3	Discussion	104
3.4	Methods	106
3.5	Supplementary information	110
4	IDENTIFYING THE INSERTION KINETICS, STABILITY AND STOICHIOMETRY OF THE BAX PORE FROM FLUORESCENCE MEASUREMENTS AND MATHEMATICAL MODELING	121
4.1	Introduction	122
4.2	Results	127
4.3	Discussion	146
4.4	Methods	149
5	CONCLUSION	157
5.1	Making models more trustworthy	158
5.2	Enlisting machines for efficient modeling	160

APPENDIX A OPTIMIZING RING ASSEMBLY REVEALS THE STRENGTH OF WEAK INTERACTIONS	163
Supplementary information	170
APPENDIX B THE NOISE IS THE SIGNAL: INFORMATION FLOW IN SINGLE CELLS AND CELLU-	
LAR POPULATIONS	237
Supplementary information	259

TO MY DAUGHTERS,
LEAH AYSUN AND EVA SELIN.

Acknowledgments

THANK YOU first and foremost to my parents, Robert Bachman and Ayla Nutku Bachman, for having given me my curiosity, ambition, and values of education and service. To my father: that in some ways this work reflects your own longstanding inquiries makes it all the more rewarding for me. Thank you for passing on this journey. To my mother: I cannot hear your words, but I can see you smiling.

Thank you to my uncle, the late physicist Yavuz Nutku, whose passion for and dedication to science—as a way of life and as a means of contributing to the world—were a major inspiration in my decision to become a scientist.

Thank you to my advisor, Peter Sorger, who granted me the intellectual freedom to pursue a peculiar project that was simultaneously futuristic and old-fashioned. Thank you for creating a fantastic environment, letting me find my own way, and giving me support and guidance when I needed it.

To my thesis and dissertation advisory committee members, Tim Mitchison, Tony Letai, Jeremy Gunawardena, and Jagesh Shah: thank you for your time, attention, and guidance over the years. It has been a pleasure and an honor to share this work with you.

Thank you to my mentors, Walter Fontana and Jeremy Gunawardena. Walter's vision for systems biology brought me into the field from computer science, and I have enjoyed the many hours we have spent at the whiteboard "seeing the same ghosts." Jeremy deserves special thanks for having gone out of his way to encourage me before I had applied to graduate school, when I was an outsider with no biology background. Taking and then twice teaching SB200 with Jeremy was a rich learning experience for which I am thankful.

I have been fortunate to have had great collaborators who became mentors and ultimately great friends. Eric Deeds, coauthor of two chapters in this thesis, taught me thermodynamics, kinetics, and a theorist's way of looking at problems. Jeremy Muhlich, coauthor of **PySB**, tremendously advanced my understanding of software design and development and helped me to find bugs after I'd gotten stumped. My baymate Josh Sims taught me all of the molecular biology and biochemistry I know and spawned my interest in the Bcl-2 family. In addition to assisting me in most of the experiments in this thesis—from making buffers to fixing the HPLC—Josh entertained hours of speculation and discussion about the Bcl-2 family, science, and life. Thank you for your generosity, integrity, and friendship.

Thank you to my coauthors and collaborators, in particular Justin Kale, with whom I spent countless hours on Skype discussing the finer points of Bax biochemistry—and it was always enjoyable. Thank you for being such a skilled, thoughtful, and open-minded partner on these projects. Thank you to David Andrews for pursuing and supporting the collaboration; to Ryan Suderman for sharing a great idea and

doing great work; to Carlos Lopez for his enthusiasm and for evangelizing our approach; to Thomas Graham and Nate Lord for the microfluidics and microscopy sessions; and to Ben Gyori for sharing both conceptual insight and practical advice.

To my '08 graduate school classmates: thank you for making it all a lot more fun. Thank you to all the members of the Sorger lab for being such a brilliant, fun, and generous group.

To Sam Reed, Chris Bird, Becky Ward, Lab Ops, and the rest of the staff of the Harvard Systems Biology Department: you have made this a wonderful place to work and study.

Thank you to my sister Nurjana for supporting me on this path and for always modeling excellence; to the Bachman and Farago families for all of the support and encouragement over the years; and to my daughters, Leah and Eva, for joy. For “is alysis your friend?” and “big mechazoom!” and so many other moments of levity and perspective.

Finally, thank you to my wife, Anna Farago, without whom this would not have been possible. Anna's own achievements as a scientist and physician-in-training were an inspiration for me when we met and I was considering applying to graduate school. Since then she has been my partner, coach, editor, and friend. Thank you for all your love and support. I owe this to you.

The language used by biologists for verbal communications...is not unlike that used by stock market analysts. Both are vague (e.g., “a balance between pro- and anti-apoptotic Bcl-2 proteins appears to control the cell viability, and seems to correlate in the long term with the ability to form tumors”) and avoid clear predictions.

Yuri Lazebnik, “Can a Biologist Fix a Radio?”
(Lazebnik, 2002)

1

Introduction

Since the molecular biology revolution of the mid-20th century, biology has concerned itself largely with the discovery of molecular mechanisms (Craver and Darden, 2013). And for good reason: in addition to satisfying our fundamental curiosity about how biological systems work, mechanistic knowledge has important practical benefits. Better understanding of molecular mechanisms can improve the efficiency of drug discovery and help identify patients likely to benefit from particular therapies (Insel et al., 2015). Mechanistic knowledge about drugs and their targets is more transferable across disease contexts than empirical observations from a single context; this has already been observed in cancer, where the use of drugs targeted to specific genomic alterations has in some cases proven successful for cancers originating from different tissues (Yan et al., 2014). Finally, a mechanistic rationale for therapy is satisfying for decision-makers, giving increased confidence to physicians, patients and regulators alike (“Mechanism matters.” 2010).

The development of modern genetic and biochemical methods ushered in tremendous growth in knowledge about biological mechanisms, culminating in a series of seemingly exhaustive studies of the

identities of cellular components and their functions (Lander et al., 2001; Venter et al., 2001; Cancer Genome Atlas Research Network et al., 2013; ENCODE Project Consortium, 2012). This knowledge has not been limited to the identification of components: databases such as Pathway Commons enumerate over a million molecular interactions, grouped into more than 30,000 pathways (Cerami et al., 2011). And yet, despite this glut of information about biological mechanisms, it remains difficult to explain deceptively simple empirical observations in mechanistic terms: for example, what makes one cell type sensitive to an anti-cancer drug and another not? The problem is not that we have no answers to this kind of question, but far too many possibilities—with the precise mechanism highly context-dependent and usually difficult to discern *a priori*.

In his classic 2002 essay, “Can a Biologist Fix a Radio?—Or, what I learned while studying apoptosis,” Yuri Lazebnik described how a subfield of biology could paradoxically become more confusing as it matured, rather than less:

At some point...the field reaches a stage at which models, that seemed so complete, fall apart, predictions that were considered so obvious are found to be wrong, and attempts to develop wonder drugs largely fail. This stage is characterized by a sense of frustration at the complexity of the process, and by a sinking feeling that despite all that intense digging the promised cure-all may not materialize. In other words, the field hits the wall, even though the intensity of research remains unabated for a while, resulting in thousands of publications, many of which are contradictory or largely descriptive....This stage can be summarized by the paradox that the more facts we learn the less we understand the process we study. (Lazebnik, 2002)

According to this view, what makes modern biology challenging is not that so much remains unknown, but that we have a limited ability to integrate what we have already discovered.

In clinical applications, recent experience with targeted therapies in cancer provide a powerful demonstration of the tremendous potential of applying mechanistic knowledge to drug discovery—as well as the limitations of current approaches. On the one hand, the BRAF-specific kinase inhibitor vemurafenib has yielded remarkably dramatic clinical responses in melanoma patients with BRAF V600E mutant tumors (Chapman et al., 2011). On the other hand, the drug has proved surprisingly ineffective in colorectal cancers harboring the same mutation (Roth et al., 2010). This clinical finding was later attributed to feed-

back regulation by EGFR, a receptor that is highly expressed in colorectal carcinoma but not melanoma cells (Prahallad et al., 2012).

In this and similar examples, the efficacy of a “magic bullet” drug aimed at a specific genetic alteration was found to depend strongly on the cellular context. Indeed, context-dependence of this type is the rule rather than the exception in biological systems, making robust mechanistic explanations and predictable design difficult (Shrager, 2003; Banerji, 2013; Del Vecchio, 2015). Given that EGFR and BRAF are among the most-studied genes in molecular biology, the appearance of unanticipated interactions among their alterations reflects the challenge of integrating knowledge about mechanism into a predictive framework. This example involves a *single* targeted drug; as the medical community turns its attention to the use of combinations of *multiple* targeted agents, the need for a more robust framework for reasoning about mechanisms will become all the more acute (Maione et al., 2006).

1.1 MATHEMATICAL MODELS INTEGRATE KNOWLEDGE ABOUT MECHANISMS

The need for a more integrated approach to understanding the behavior of biological systems was one reason for the emergence of the field of systems biology, around the time of Lazebnik’s essay. From the beginning, it was envisioned that this new field would have both experimental and theoretical components, focusing on quantitative phenomena and explanations (Kirschner, 2005; Wingreen and Botstein, 2006; Liu, 2005). Quantitative models, which already had a long and successful—if underappreciated—track record in biology (Gunawardena, 2014), were seen as a key tool for reasoning about complex biological systems (Sorger, 2005; Tyson, 2007).

Over the last decade, quantitative modeling has shown promise as a tool for explaining and predicting the behavior of biological systems. Models have been used to explore a number of complex processes, including the cell cycle of budding yeast (K. C. Chen et al., 2000), mammalian growth factor signaling (W. W. Chen et al., 2009; Blinov et al., 2006), receptor-mediated apoptosis (Albeck et al., 2008a; Bentele et al., 2004; Lindner et al., 2013), NF- κ B signaling (Cheong et al., 2008), the microbe *M. genitalium* (Karr et al., 2012), and human metabolism (Thiele et al., 2013).

As it is generally practiced, modeling in systems biology involves enumerating a set of assumptions

about a subset of cellular components and their interactions, drawn from the literature; model scope is itself an assumption and remains largely a matter of intuition. Model assumptions are encoded into a formal representation allowing inferences or predictions to be made, which are then compared to data or analyzed mathematically.

There are three principal difficulties arising from the current approach. First, by encoding assumptions directly into a low-level modeling formalism (such as equations), models tend to mirror (or even magnify) the complexity of the systems that they describe. As a model’s biological scope increases, its structure and assumptions become increasingly obscure, making it difficult to reuse or extend. This is a primary barrier to models becoming reusable community resources. Second, the use of a single model results in brittle explanations, offering little insight into the connection between a model’s assumptions and its behavior. The uniqueness of the model in explaining the data is unknown, and the model may be “right” (fit the data) for the wrong reasons. Finally, model construction remains a highly laborious process, requiring significant amounts of time and expertise (this point is revisited in [Chapter 5](#)).

As a consequence of these persistent difficulties with current modeling approaches, the particular requirement Lazebnik identified—of a durable framework for integrating mechanistic explanations of complex phenomena—remains unsatisfied. The need for more robust explanatory and predictive frameworks is all the more apparent considering the growth in large, complex datasets (e.g., [Vidović et al., 2014](#)). Even in this age of ‘omics, explanations of complex biological phenomena tend to focus on cherry-picked examples, are local to particular experiments and datasets, are only informally grounded in prior knowledge of mechanisms, and as discussed above, are rarely robust across biological contexts.

1.2 PROGRAMMATIC ENSEMBLE MODELING

In this thesis I discuss an approach, *programmatic ensemble modeling*, to address these bottlenecks in using models to interpret biological phenomena.

Programmatic modeling refers to the practice of using tools and approaches from computer software to build models that are transparent, reusable, and extensible. Despite their complexity, biological models contain many repeated elements, such as low-level biophysical processes that are highly stereotyped

(binding, post-translational modification, transport, assembly, etc.). The use of abstraction allows these common elements to be represented by reusable functions, masking the low-level details and reducing implementation errors. Models can be hierarchically decomposed into modules at varying scales, from self-contained pathways to individual mechanistic elements that are common to many model variants (Figure 1.1). This approach was explored in the LISP framework *little b* (Mallavarapu et al., 2009), and extended in PySB, described in Chapter 2 (Lopez et al., 2013).




Figure 1.1: Programmatic modeling allows hierarchical decomposition of complex biological processes. **(A)** Alternative models can be built by recombining reusable elements at multiple scales, termed “modules” (larger scale) or “motifs” (smaller scale). At the lowest level, “macros” implement abstractions of common biochemical processes. **(B)** In conventional modeling practice, a model is encoded directly in a mathematical representation, masking the many intermediate assumptions. (Figure adapted from Chapter 2.)

Programmatic construction of models is essential for building and analyzing multiple model variants, as it makes the mechanistic differences between alternative models explicit. For example, two models of the extrinsic apoptosis pathway, EARM 2.0-M2a and EARM 2.0-M3a, share a common implementation of the upstream reactions (*rec_to_bid*; Figure 1.1), but differ in the modules used to implement the mitochondrial reactions (*direct* vs. *indirect*). The high-level differences between these two modules can then be seen immediately in terms of the different mechanistic motifs they contain. At each level of the hierarchy, differences between model alternatives are captured in terms of a high-level vocabulary that maps closely to the terms used by biologists.

Ensemble modeling involves formulating a collection of alternative models of a process and then evaluating them against data. After model calibration (i.e., estimation of unknown parameters), models with poor fit can be discarded using statistical techniques for model discrimination; predictions of the remain-




Figure 1.2: Single vs. ensemble modeling. **(A)** In conventional modeling practice, a single model is formulated and fit to data, leading to iterative refinements until a satisfactory fit is obtained. **(B)** In ensemble modeling, many candidate models are defined *a priori* based on prior knowledge; these models are then ranked or grouped based on their fit to data. The ensemble of models can then be expanded to incorporate new kinds of assumptions or used for experimental design or prediction.

ing models with good fit can be pooled as a way of managing the uncertainty about which model is correct (Figure 1.2) (Burnham and Anderson, 2002).

Biology is a good candidate for multi-model analysis because biochemistry provides a solid theoretical foundation for the mechanistic building blocks (Burnham and Anderson, 2002). Though not the predominant approach in systems biology, a number of previous studies have used discrimination among model variants to refine hypotheses about biological mechanisms (Kleiman et al., 2011; Muzzey et al., 2009; Albeck et al., 2008b). Ensemble modeling was first explicitly proposed as a modeling method for biology by Stelling and colleagues, chiefly for its value in managing uncertainty about mechanisms in signaling (Kuepfer et al., 2007). A tool for model enumeration from a single master model, modelMaGe, has also been described (Schaber et al., 2011).

Ensemble modeling offers several advantages over the practice of building and refining single models (Figure 1.2):

- *Identifies when simple models are sufficient to fit the data.* Inclusion of both simple and complex

models in an ensemble highlights which mechanisms are primarily responsible for the experimental observations, yielding insight into the key drivers of system behavior. Conversely, evaluation of the model ensemble across multiple datasets can also identify limitations on the scope of the simpler models.

- *Groups models having equivalent behaviors, allowing insight into mechanisms.* Particular combinations of mechanistic assumptions can lead to shared behaviors among models that may not *a priori* be seen as similar. The process of model calibration and discrimination yields a classification of models into groups that succeed or fail in similar ways, providing insight into the roles of specific combinations of assumptions in driving model behavior.
- *Highlights ambiguities in mechanistic explanations.* It is commonly appreciated that models (both qualitative and quantitative) offered as explanations of biological phenomena are not unique, in that other explanations could prove equally valid. The model ensemble provides a quantitative framework for assessing this type of ambiguity.
- *Facilitates experimental design for model refinement.* A subset of models that are equally valid given existing experimental data may be able to be discriminated with carefully chosen experiments. Predictions of the model ensemble under different experimental conditions can be used to design experiments likely to result in clarification of model ambiguity.
- *Manages uncertainty through ensemble predictions.* Where model predictions, rather than model structure, are of primary interest (e.g., in a precision medicine context), predictions from the ensemble of validated models can be pooled to account for uncertainty in the underlying mechanism; this practice is widely used in weather forecasting (Leutbecher and Palmer, 2008).

Because of these benefits, programmatic enumeration of multiple mechanistic models has the potential to facilitate the discovery and refinement of mechanistic descriptions of biological processes. In this thesis I apply this approach to mechanistic questions in apoptotic pore formation, reviewed briefly in the following section.

1.3 MOLECULAR MECHANISMS IN APOPTOSIS

Apoptosis is a programmed cell death pathway that eliminates cells as part of normal cellular turnover and in response to stress and damage (Green, 2011). It is marked by several morphologic features, including membrane blebbing and nuclear fragmentation, and is defined biochemically by the activation of a family of cysteine proteases known as caspases (Galluzzi et al., 2014). Once activated by proteolytic cleavage, the subgroup of caspases known as “executioners” cleave a large number of downstream substrates, leading to the breakdown of the cell.

A key step in the activation of executioner caspases is mitochondrial outer membrane permeabilization, or MOMP (Tait and Green, 2010). In MOMP, the formation of pores in the mitochondrial outer membrane results in the release of a number of pro-apoptotic factors from the mitochondrial intermembrane space into the cytosol, where they trigger the activation of executioner caspases (Figure 1.3). In most cases the post-MOMP reactions are rapid and complete, making MOMP the “point of no return” (Goldstein et al., 2000; Rehm et al., 2002; Albeck et al., 2008a). MOMP is an essential process for stress-mediated (intrinsic) apoptosis and, in many cell types, for receptor-mediated (extrinsic) apoptosis as well (Barnhart et al., 2003; Maas et al., 2013). Because of its role in governing life-death decisions, the tendency of a cell’s mitochondria to undergo MOMP is highly predictive of cell death in response to anti-cancer drugs and, ultimately, clinical responses (Chonghaile et al., 2011; Vo et al., 2012).




Figure 1.3: General mechanism of intrinsic apoptosis. Cellular stress or other damage leads to the activation of one or more BH3-only proteins, which converge on the mitochondrial outer membrane to trigger MOMP by interactions with other Bcl-2 proteins. MOMP leads to the release of pro-apoptotic proteins such as cytochrome c and Smac, which activation executioner caspases, triggering the breakdown of the cell.

MOMP is regulated by an evolutionarily-related family of proteins known as the Bcl-2 family, which contains both pro- and anti-apoptotic members (Figure 1.4) (Chipuk et al., 2010; Youle and Strasser, 2008; Czabotar et al., 2014). The pro-apoptotic “effectors,” Bax and Bak, undergo an activating conformational change allowing them to oligomerize and form the pores in the mitochondrial outer membrane (Hsu et al., 1997; Hsu and Youle, 1998; Hsu and Youle, 1997). This conformational change is triggered by binding of members of the BH3-only subfamily of Bcl-2 proteins, which share only one of the four Bcl-2 homology (BH) regions with the other members (Llambi et al., 2011; Kim et al., 2009). The anti-apoptotic family members, which include the known oncogenes Bcl-2, Bcl-XL, and Mcl-1, prevent death by binding and inhibiting both the BH3-only activators (Certo et al., 2006; Letai et al., 2002; Kim et al., 2006) and the activated effectors (L. Chen et al., 2005; Willis et al., 2007; Willis et al., 2005). The additional subgroup of BH3-only proteins known as “sensitizers,” though not able to trigger activation of Bax and Bak, bias the cell to pro-death stimuli by inhibiting the anti-apoptotic proteins (Letai et al., 2002). The dual role of activator BH3-only proteins in activating Bax and Bak—directly, by triggering activating conformational changes, as well as indirectly, by inhibiting the anti-apoptotics—was the basis of a longstanding debate over which of these two modes was of greater importance, a topic I revisit in Chapter 2 (Leber et al., 2007; Leber et al., 2010; Chipuk et al., 2010).




Figure 1.4: Interactions among Bcl-2 family members. Binding of the BH3-only activators triggers an activating conformational change in the effectors Bax and Bak (indicated by *). The anti-apoptotic members inhibit this process in two ways: first, by binding the activator BH3-onlies, they prevent activation of Bax and Bak. Second, by binding specifically to the activated forms of Bax and Bak they prevent downstream formation of pores. The sensitizer BH3-only proteins are unable to activate Bax and Bak, but by binding to the anti-apoptotics they enable both activation and pore formation. Recent evidence suggests that Bok, which shares structural characteristics with Bax and Bak, can form pores but does not depend on BH3-only proteins for activation.

Several Bcl-2 proteins have been crystallized in their soluble truncated forms, aiding in the understanding of the structural basis of their interactions (Ku et al., 2011; Suzuki et al., 2000; Aritomi et al., 1997; Day

et al., 2005). The anti-apoptotics appear to share a common mechanism for binding to both BH₃-only proteins and effectors, in which a conserved hydrophobic groove spanning the BH₃ and BH₁ regions of the anti-apoptotic protein binds tightly to the BH₃ region of the effector or BH₃-only protein, inhibiting its activity. A recent study has also suggested that an interaction of the BH₄ domain of Bcl-2 with a noncanonical pocket on Bax may also have a role in inhibition of pore formation (Barclay et al., 2015).

The effectors Bax and Bak are subject to a complex series of regulatory steps en route to activation, oligomerization, and pore formation; the kinetics and structural basis of these steps are the focus of Chapter 3. While Bak is constitutively bound to the mitochondrial membrane, Bax is predominantly cytosolic and undergoes an activating conformational change associated with translocation to mitochondrial membranes and exposure of the N-terminus (Hsu et al., 1997; Hsu and Youle, 1998; Hsu and Youle, 1997). The NMR structure of Bax revealed that in its soluble conformation, Bax's hydrophobic C-terminal helix is folded into its hydrophobic groove, preventing membrane insertion (Suzuki et al., 2000). Binding of a chemically stabilized BH₃-only peptide to a "trigger site" on a pocket of Bax opposite the hydrophobic groove triggers the release of the C-terminal helix and the exposure of the α_1 and α_2 helices (Gavathiotis et al., 2008; Gavathiotis et al., 2010). *In silico* screening for ligands binding this trigger site led to the development of a small molecule Bax activator, suggesting its relevance for activation (Gavathiotis et al., 2012). Structural analysis has suggested that the conformational change is propagated through the core of Bax through a network of charge reorganizations (Ionescu et al., 2012).

There is also strong evidence that binding of BH₃-only proteins directly to the hydrophobic groove of both Bax and Bak plays an important role in their activation (Czabotar et al., 2013; Dai et al., 2011; Moldoveanu et al., 2013). A crystal structure of the BH₃-only peptide of Bid in complex with the hydrophobic groove of Bax not only suggested that these proteins bind in a similar fashion as BH₃-onlies and anti-apoptotics, but also provided a structural explanation for why some BH₃-only proteins can activate Bax and Bak and others not (Czabotar et al., 2013).

The relative roles of the rear pocket trigger site and the hydrophobic groove in Bax activation is currently unclear. One possibility is that binding at the rear pocket is responsible for the initial release of the N- and C-terminal helices but is not sufficient for full activation, with this latter step mediated by

binding of the hydrophobic groove (Czabotar et al., 2014). However, this interpretation raises additional questions such as whether these two distinct sites could be engaged by a BH₃-only protein in a single sequential interaction, or whether Bax dissociates from its activator after the first step and is subsequently rebound by a (potentially different) activator.

The structure and stoichiometry of the assembled apoptotic pore has been a key subject of investigation, with important aspects remaining unresolved, including whether the stoichiometry is fixed or variable (Volkman et al., 2013; Borner and Andrews, 2014). A patch-clamp study using isolated mitochondria suggested that the Bax and Bak pore stoichiometry is fixed at roughly nine monomers consistent with a barrel-stave configuration (Martinez-Caballero et al., 2009). Kinetic analysis has been used to suggest that the minimal stoichiometry of the Bax pore is between two and four (Saito et al., 2000; Schlesinger and Saito, 2006), though this analysis was based on an interpretation of the Hill coefficient as a stoichiometry, which is not generally valid (Weiss, 1997). A time-lapse study of Bax-mediated dye release from giant unilamellar vesicles indicated that smaller molecules were released earlier than large ones, suggesting a growing pore (Bleicken et al., 2013). Single-molecule microscopy of the formation of Bax oligomers on planar lipid bilayers indicated that Bax oligomers are variable in size but consist of dimeric units, consistent with previous crosslinking studies (Subburaj et al., 2015; Kim et al., 2006).

The question of stoichiometry is related to the question of pore stability: if pore formation involves labile assemblies of Bax monomers or dimers, this might suggest that pores could form transiently and subsequently dissipate, providing another point of regulation. A study measuring the dye release profile of outer membrane vesicles prepared from isolated mitochondria suggested that pores could indeed collapse (Gillies et al., 2015); however, studies using the $\alpha 5$ peptide from Bax's pore forming region, which is sufficient to form pores, showed that once formed, Bax pores tend to remain open (Fuertes et al., 2010). Additional studies will be required to fully resolve these questions.

Substantial evidence now indicates that dimerization of both Bax and Bak is mediated by symmetric BH₃:groove interactions (Dewson et al., 2008; Dewson et al., 2012), rather than in an asymmetric "head-to-tail" fashion (Bogner et al., 2010). Mounting evidence that activation of Bax and Bak by BH₃-only proteins and Bax/Bak homodimerization are *both* mediated by canonical BH₃:groove interactions presents

a conundrum: how do the BH₃-only activators engage the Bax/Bak groove to trigger a conformational change without preventing subsequent homodimerization and pore formation? One explanation is that engagement of the Bax/Bak groove by BH₃-onlies results in a conformational change in the groove itself that is less favorable for the BH₃-onlies, causing their rapid dissociation in a “hit-and-run” fashion. However, it has also been shown by FRET that Bid and Bax remain bound at equilibrium, when most Bax is in pores (Lovell et al., 2008). Co-immunoprecipitation of activated Bax using the 6A7 antibody pulls down tBid, another indication that Bid and Bax remain bound (Kim et al., 2009). This suggests a dynamic equilibrium with rapid exchange between BH₃:Bax and Bax:Bax dimers, or a noncanonical interaction between BH₃-onlies and activated Bax at a site other than its hydrophobic groove.

Studies with Bax-derived peptides have shown that Bax’s own BH₃ region can trigger activation of full-length Bax, suggesting positive feedback by auto-activation (Tan et al., 2006). Theoretical analysis has suggested that Bax autoactivation could play an important role in making pore formation bistable and thereby irreversible (Cui et al., 2008). Experiments using purified Bax activated by heat rather than BH₃-only proteins showed that while Bax mutated at the trigger site (K21E) was inactive, the double mutant with the complementary mutation in Bax’s own BH₃ region was able to form pores, suggesting a role for Bax auto-activation *in trans* via the trigger site (Gavathiotis et al., 2010). However, because dimerization of Bax via a symmetric BH₃:groove interaction would sequester the BH₃ region and thereby prevent auto-activation, it is unclear how significant this type of interaction is in typical scenarios.

1.4 SYSTEMS BIOCHEMISTRY ALLOWS QUANTITATIVE CHARACTERIZATION OF MECHANISMS

Kinetic analysis of protein interactions *in vitro* has been a mainstay of biochemistry since the identification of the enzyme-substrate complex by Michaelis and Menten (Gunawardena, 2012). However, modern systems biology has focused on using dynamical models either to derive theoretical conclusions or to interpret experimental data from whole cells. While these studies are valuable, it is frequently difficult to make the necessary (often complex) cellular perturbations to validate model conclusions. For example, while Albeck et al. demonstrated that MOMP was responsible for the snap-action kinetics of caspase activation in single cells, subsequent biochemical studies suggested limited support for the proposed

mechanism (Albeck et al., 2008b).

Reconstitution biochemistry involving complex mixtures of proteins, membranes and cellular extracts offers a middle ground between binary protein-protein interaction studies and cell-based experiments. This approach, which has also been termed “meso-scale systems biology” (Kirschner, 2005) or “synthetic biology of minimal systems,” (Schwille and Diez, 2009) provides a fertile ground for detailed and quantitative characterization of mechanisms. In the *in vitro* setting, control over absolute concentrations of individual elements allows for systematic experimental perturbations and mechanistic interpretation. Because *in vitro* studies reduce many of the experimental barriers that hold back many modeling studies, they also allow for deeper exploration of mechanistic alternatives and highlight the strengths and weaknesses of modeling methods.

In vitro studies of Bcl-2 family interactions using synthetic liposomes or isolated mitochondria have provided many important insights into pore formation mechanisms. A key step was the identification of the minimal set of components required to reconstitute pore formation activity *in vitro*, including an early analysis of the importance of the particular lipid constituents of the membrane in facilitating Bax activation (Kuwana et al., 2002). Many subsequent reconstitution experiments established the identity of activators (Du et al., 2011), the sequence of steps involved in Bax activation (Lovell et al., 2008), and the role of specific lipid and protein constituents in the pore formation process (Schafer et al., 2009; Shamas-Din et al., 2015; Shamas-Din et al., 2013).

1.5 SYNOPSIS OF THESIS

In Chapter 2 of this thesis, I introduce PySB, a framework for programmatic modeling of biochemical systems. Describing biochemical systems using tools and approaches from computer programming allows complex pathways to be described with a high-level vocabulary that is a good conceptual fit with the language naturally used by biologists. Importantly, this enables alternative models of a process to be easily and transparently enumerated, allowing for calibration and discrimination of model ensembles. To demonstrate this approach, I formulate several alternative models of Bcl-2 mediated regulation of MOMP; these include several previously published in the literature as well as newly formulated “di-

rect,” “indirect,” and “embedded together” models. All models are described using reusable elements at multiple scales, termed *modules*, *motifs* and *macros*, making differences between the models immediately apparent upon inspection of the model code.

In [Chapter 3](#), I use an ensemble of kinetic models to identify a novel feature of the Bax pore formation pathway, the formation of a structurally distinct intermediate preceding pore formation. Discrimination among an ensemble of models representing differing numbers of conformational states suggests that Bax adopts three or four distinct conformational states en route to pore formation. The model allows data from labeled Bax variants with widely differing activities to be integrated into a single coherent picture of the activation sequence. Analysis of the three conformation model shows that the first transition is fast and is associated with binding of BH₃-only activators and a concurrent increase in hydrophobicity of the N-terminal and BH₃ regions of Bax. The second transition is slower and is associated with Bax oligomerization and pore formation. The hydrophobicity increase of the C-terminal region occurs on a slower timescale coordinated with insertion of pore-forming helices, suggesting that this is a later step than previously thought. Analysis of Bax point mutants identified through cancer genome sequencing results in the identification of mutants that have blocks in pore formation at the intermediate step, indicating that the two transitions are structurally independent.

In [Chapter 4](#) I address the closely related question of how concentrations of the various constituents of a minimal pore formation system—cBid, Bax, and synthetic liposomes—affect the dynamics of Bax insertion and permeabilization. By systematically exploring the concentration phase space, I show that many mechanistic details of the various types of interactions—Bid:lipid, Bax:lipid, Bid:Bax, and Bax:Bax—must be accounted for in order to obtain a complete and consistent explanation of the data. In addition, overall permeabilization kinetics is, paradoxically, slowed down by the nearly irreversible recruitment of Bax to existing pores. I find that pores are stable rather than transient, and consist minimally of four subunits.

Finally, in [Chapter 5](#), I discuss the outlook for extending the modeling approaches described in this thesis to larger and more complex systems. Adoption of a “iterate modeling” methodology will likely make construction of large models tractable for a larger community of modelers. In addition, recent

developments in machine reading of the scientific literature may help with the assembly and translation of prior knowledge into executable models, greatly reducing the amount of time and effort involved.

References

- Albeck, J. G., J. M. Burke, B. B. Aldridge, M. Zhang, D. A. Lauffenburger, and P. K. Sorger (2008a). “Quantitative analysis of pathways controlling extrinsic apoptosis in single cells”. *Molecular Cell* 30.1 (Apr. 2008), pp. 11–25.
- Albeck, J. G., J. M. Burke, S. L. Spencer, D. A. Lauffenburger, and P. K. Sorger (2008b). “Modeling a snap-action, variable-delay switch controlling extrinsic cell death”. *PLoS Biology* 6.12 (Dec. 2008), pp. 2831–2852.
- Aritomi, M., N. Kunishima, N. Inohara, Y. Ishibashi, S. Ohta, and K. Morikawa (1997). “Crystal structure of rat Bcl-xL. Implications for the function of the Bcl-2 protein family.” *The Journal of biological chemistry* 272.44 (Oct. 1997), pp. 27886–27892.
- Banerji, A. (2013). “An attempt to construct a (general) mathematical framework to model biological ”context-dependence”.” *Systems and Synthetic Biology* 7.4 (Dec. 2013), pp. 221–227.
- Barclay, L. A. et al. (2015). “Inhibition of Pro-Apoptotic BAX by a Noncanonical Interaction Mechanism.” *Multiple values selected* 57.5 (Feb. 2015), pp. 873–886.
- Barnhart, B. C., E. C. Alappat, and M. E. Peter (2003). “The CD95 type I/type II model”. *Seminars in Immunology* 15.3 (June 2003), pp. 185–193.
- Bentele, M., I. Lavrik, M. Ulrich, S. Stösser, D. W. Heermann, H. Kalthoff, P. H. Krammer, and R. Eils (2004). “Mathematical modeling reveals threshold mechanism in CD95-induced apoptosis”. *The Journal of Cell Biology* 166.6 (Sept. 2004), pp. 839–851.
- Bleicken, S., O. Landeta, A. Landajuela, G. Basañez, and A. J. García-Sáez (2013). “Proapoptotic Bax and Bak proteins form stable protein-permeable pores of tunable size.” *Journal of Biological Chemistry* 288.46 (Nov. 2013), pp. 33241–33252.
- Blinov, M. L., J. R. Faeder, B. Goldstein, and W. S. Hlavacek (2006). “A network model of early events in epidermal growth factor receptor signaling that accounts for combinatorial complexity.” *Bio Systems* 83.2-3 (Jan. 2006), pp. 136–151.
- Bogner, C., B. Leber, and D. W. Andrews (2010). “Apoptosis: embedded in membranes.” *Current Opinion in Cell Biology* 22.6 (Dec. 2010), pp. 845–851.
- Borner, C. and D. W. Andrews (2014). “The apoptotic pore on mitochondria: are we breaking through or still stuck?” *Cell Death and Differentiation* 21.2 (Feb. 2014), pp. 187–191.
- Burnham, K. P. and D. R. Anderson (2002). *Model selection and multi-model inference: a practical information-theoretic approach*. 2nd. New York: Springer.

- Cancer Genome Atlas Research Network et al. (2013). “The Cancer Genome Atlas Pan-Cancer analysis project.” *Nature Genetics* 45.10 (Oct. 2013), pp. 1113–1120.
- Cerami, E. G., B. E. Gross, E. Demir, I. Rodchenkov, Ö. Babur, N. Anwar, N. Schultz, G. D. Bader, and C. Sander (2011). “Pathway Commons, a web resource for biological pathway data.” *Nucleic Acids Research* 39.Database issue (Jan. 2011), pp. D685–90.
- Certo, M., V. Del Gaizo Moore, M. Nishino, G. Wei, S. Korsmeyer, S. A. Armstrong, and A. Letai (2006). “Mitochondria primed by death signals determine cellular addiction to antiapoptotic BCL-2 family members”. *Cancer Cell* 9.5 (May 2006), pp. 351–365.
- Chapman, P. B. et al. (2011). “Improved survival with vemurafenib in melanoma with BRAF V600E mutation.” *The New England journal of medicine* 364.26 (June 2011), pp. 2507–2516.
- Chen, K. C., A. Csikasz-Nagy, B. Györfy, J. Val, B. Novak, and J. J. Tyson (2000). “Kinetic analysis of a molecular model of the budding yeast cell cycle.” *Molecular Biology of the Cell* 11.1 (Jan. 2000), pp. 369–391.
- Chen, L. et al. (2005). “Differential Targeting of Prosurvival Bcl-2 Proteins by Their BH₃-Only Ligands Allows Complementary Apoptotic Function”. *Molecular Cell* 17.3 (Apr. 2005), pp. 393–403.
- Chen, W. W., B. Schoeberl, P. J. Jasper, M. Niepel, U. B. Nielsen, D. A. Lauffenburger, and P. K. Sorger (2009). “Input-output behavior of ErbB signaling pathways as revealed by a mass action model trained against dynamic data”. *Molecular Systems Biology* 5, p. 239.
- Cheong, R., A. Hoffmann, and A. Levchenko (2008). “Understanding NF-kappaB signaling via mathematical modeling.” *Molecular Systems Biology* 4.1, p. 192.
- Chipuk, J. E., T. Moldoveanu, F. Llambi, M. J. Parsons, and D. R. Green (2010). “The BCL-2 family reunion.” *Molecular Cell* 37.3 (Feb. 2010), pp. 299–310.
- Chonghaile, T. N. et al. (2011). “Pretreatment Mitochondrial Priming Correlates with Clinical Response to Cytotoxic Chemotherapy”. *Science (New York, NY)* 334.6059 (Nov. 2011), pp. 1129–1133.
- Craver, C. F. and L. Darden (2013). *In Search of Biological Mechanisms: Discoveries across the Life Sciences*. Chicago, IL: University of Chicago Press.
- Cui, J., C. Chen, H. Lu, T. Sun, and P. Shen (2008). “Two independent positive feedbacks and bistability in the Bcl-2 apoptotic switch”. *PLoS ONE* 3.1, e1469.
- Czabotar, P. E., G. Lessene, A. Strasser, and J. M. Adams (2014). “Control of apoptosis by the BCL-2 protein family: implications for physiology and therapy.” *Nature Reviews Molecular Cell Biology* 15.1 (Jan. 2014), pp. 49–63.
- Czabotar, P. E. et al. (2013). “Bax crystal structures reveal how BH₃ domains activate Bax and nucleate its oligomerization to induce apoptosis.” *Cell* 152.3 (Jan. 2013), pp. 519–531.

- Dai, H., A. Smith, X. W. Meng, P. A. Schneider, Y.-P. Pang, and S. H. Kaufmann (2011). “Transient binding of an activator BH₃ domain to the Bak BH₃-binding groove initiates Bak oligomerization.” *The Journal of Cell Biology* 194.1 (July 2011), pp. 39–48.
- Day, C. L., L. Chen, S. J. Richardson, P. J. Harrison, D. C. S. Huang, and M. G. Hinds (2005). “Solution structure of prosurvival Mcl-1 and characterization of its binding by proapoptotic BH₃-only ligands.” *The Journal of biological chemistry* 280.6 (Feb. 2005), pp. 4738–4744.
- Del Vecchio, D. (2015). “Modularity, context-dependence, and insulation in engineered biological circuits.” *Trends in biotechnology* 33.2 (Feb. 2015), pp. 111–119.
- Dewson, G., S. Ma, P. Frederick, C. Hockings, I. Tan, T. Kratina, and R. M. Kluck (2012). “Bax dimerizes via a symmetric BH₃:groove interface during apoptosis.” *Cell Death and Differentiation* 19.4 (Apr. 2012), pp. 661–670.
- Dewson, G., T. Kratina, H. W. Sim, H. Puthalakath, J. M. Adams, P. M. Colman, and R. M. Kluck (2008). “To trigger apoptosis, Bak exposes its BH₃ domain and homodimerizes via BH₃:groove interactions.” *Molecular Cell* 30.3 (May 2008), pp. 369–380.
- Du, H., J. Wolf, B. Schafer, T. Moldoveanu, J. E. Chipuk, and T. Kuwana (2011). “BH₃ domains other than Bim and Bid can directly activate Bax/Bak.” *Journal of Biological Chemistry* 286.1 (Jan. 2011), pp. 491–501.
- ENCODE Project Consortium (2012). “An integrated encyclopedia of DNA elements in the human genome.” *Nature* 489.7414 (Sept. 2012), pp. 57–74.
- Fuertes, G., A. J. García-Sáez, S. Esteban-Martín, D. Giménez, O. L. Sánchez-Muñoz, P. Schwillle, and J. Salgado (2010). “Pores formed by Bax α 5 relax to a smaller size and keep at equilibrium”. *Biophysical Journal* 99.9 (Nov. 2010), pp. 2917–2925.
- Galluzzi, L. et al. (2014). “Essential versus accessory aspects of cell death: recommendations of the NCCD 2015.” *Cell Death and Differentiation* (Sept. 2014).
- Gavathiotis, E., D. E. Reyna, J. A. Bellairs, E. S. Leshchiner, and L. D. Walensky (2012). “Direct and selective small-molecule activation of proapoptotic BAX.” *Nature Chemical Biology* 8.7 (July 2012), pp. 639–645.
- Gavathiotis, E., D. E. Reyna, M. L. Davis, G. H. Bird, and L. D. Walensky (2010). “BH₃-Triggered Structural Reorganization Drives the Activation of Proapoptotic BAX”. *Molecular Cell* 40.3 (Nov. 2010), pp. 481–492.
- Gavathiotis, E. et al. (2008). “BAX activation is initiated at a novel interaction site”. *Nature* 455.7216 (Oct. 2008), pp. 1076–1081.
- Gillies, L. A., H. Du, B. Peters, C. M. Knudson, D. D. Newmeyer, and T. Kuwana (2015). “Visual and functional demonstration of growing Bax-induced pores in mitochondrial outer membranes”. *Molecular Biology of the Cell* 26.2 (Jan. 2015), pp. 339–349.

- Goldstein, J. C., N. J. Waterhouse, P. Juin, G. I. Evan, and D. R. Green (2000). "The coordinate release of cytochrome c during apoptosis is rapid, complete and kinetically invariant". *Nature Cell Biology* 2.3 (Mar. 2000), pp. 156–162.
- Green, D. R. (2011). *Means to an End: Apoptosis and Other Cell Death Mechanisms*. Cold Spring Harbor, New York: Cold Spring Harbor Laboratory Press.
- Gunawardena, J. (2012). "Some lessons about models from Michaelis and Menten." *Molecular Biology of the Cell* 23.4 (Feb. 2012), pp. 517–519.
- Gunawardena, J. (2014). "Models in biology: 'accurate descriptions of our pathetic thinking'." *BMC biology* 12.1, p. 29.
- Hsu, Y. T., K. G. Wolter, and R. J. Youle (1997). "Cytosol-to-membrane redistribution of Bax and Bcl-X(L) during apoptosis." *Proceedings of the National Academy of Sciences of the United States of America* 94.8 (Apr. 1997), pp. 3668–3672.
- Hsu, Y. T. and R. J. Youle (1997). "Nonionic detergents induce dimerization among members of the Bcl-2 family". *The Journal of biological chemistry* 272.21 (May 1997), pp. 13829–13834.
- Hsu, Y. T. and R. J. Youle (1998). "Bax in murine thymus is a soluble monomeric protein that displays differential detergent-induced conformations." *The Journal of biological chemistry* 273.17 (Apr. 1998), pp. 10777–10783.
- Insel, P. A., S. G. Amara, and T. F. Blaschke (2015). "Introduction to the theme "Precision medicine and prediction in pharmacology"." *Annual review of pharmacology and toxicology* 55, pp. 11–14.
- Ionescu, C.-M., R. Svobodová Vařeková, J. H. M. Prehn, H. J. Huber, and J. Koča (2012). "Charge profile analysis reveals that activation of pro-apoptotic regulators Bax and Bak relies on charge transfer mediated allosteric regulation." *PLoS Computational Biology* 8.6, e1002565.
- Karr, J. R., J. C. Sanghvi, D. N. Macklin, M. V. Gutschow, J. M. Jacobs, B. Bolival, N. Assad-Garcia, J. I. Glass, and M. W. Covert (2012). "A whole-cell computational model predicts phenotype from genotype." *Cell* 150.2 (July 2012), pp. 389–401.
- Kim, H., M. Rafiuddin-Shah, H.-C. Tu, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2006). "Hierarchical regulation of mitochondrion-dependent apoptosis by BCL-2 subfamilies". *Nature Cell Biology* 8.12 (Dec. 2006), pp. 1348–1358.
- Kim, H., H.-C. Tu, D. Ren, O. Takeuchi, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2009). "Stepwise activation of BAX and BAK by tBID, BIM, and PUMA initiates mitochondrial apoptosis". *Molecular Cell* 36.3 (Nov. 2009), pp. 487–499.
- Kirschner, M. (2005). "The Meaning of Systems Biology". *Cell* 121.4 (May 2005), pp. 503–504.
- Kleiman, L. B., T. Maiwald, H. Conzelmann, D. A. Lauffenburger, and P. K. Sorger (2011). "Rapid phospho-turnover by receptor tyrosine kinases impacts downstream signaling and drug binding." *Molecular Cell* 43.5 (Sept. 2011), pp. 723–737.

- Ku, B., C. Liang, J. U. Jung, and B.-H. Oh (2011). "Evidence that inhibition of BAX activation by BCL-2 involves its tight and preferential interaction with the BH₃ domain of BAX." *Cell research* 21.4 (Apr. 2011), pp. 627–641.
- Kuepfer, L., M. Peter, U. Sauer, and J. Stelling (2007). "Ensemble modeling for analysis of cell signaling dynamics." *Nature Biotechnology* 25.9 (Sept. 2007), pp. 1001–1006.
- Kuwana, T., M. R. Mackey, G. Perkins, M. H. Ellisman, M. Latterich, R. Schneider, D. R. Green, and D. D. Newmeyer (2002). "Bid, Bax, and lipids cooperate to form supramolecular openings in the outer mitochondrial membrane." *Cell* 111.3 (Nov. 2002), pp. 331–342.
- Lander, E. S. et al. (2001). "Initial sequencing and analysis of the human genome." *Nature* 409.6822 (Feb. 2001), pp. 860–921.
- Lazebnik, Y. (2002). "Can a biologist fix a radio?—Or, what I learned while studying apoptosis". *Cancer Cell* 2.3 (Sept. 2002), pp. 179–182.
- Leber, B., J. Lin, and D. W. Andrews (2010). "Still embedded together binding to membranes regulates Bcl-2 protein interactions." *Oncogene* 29.38 (Sept. 2010), pp. 5221–5230.
- Leber, B., J. Lin, and D. W. Andrews (2007). "Embedded together: the life and death consequences of interaction of the Bcl-2 family with membranes." *Apoptosis : an international journal on programmed cell death* 12.5 (May 2007), pp. 897–911.
- Letai, A., M. C. Bassik, L. D. Walensky, M. D. Sorcinelli, S. Weiler, and S. J. Korsmeyer (2002). "Distinct BH₃ domains either sensitize or activate mitochondrial apoptosis, serving as prototype cancer therapeutics". *Cancer Cell* 2.3 (Sept. 2002), pp. 183–192.
- Leutbecher, M. and T. N. Palmer (2008). "Ensemble forecasting". *Journal of Computational Physics* 227.7 (Mar. 2008), pp. 3515–3539.
- Lindner, A. U. et al. (2013). "Systems analysis of BCL2 protein family interactions establishes a model to predict responses to chemotherapy." *Cancer research* 73.2 (Jan. 2013), pp. 519–528.
- Liu, E. T. (2005). "Systems biology, integrative biology, predictive biology." *Cell* 121.4 (May 2005), pp. 505–506.
- Llambi, F., T. Moldoveanu, S. W. G. Tait, L. Bouchier-Hayes, J. Temirov, L. L. McCormick, C. P. Dillon, and D. R. Green (2011). "A unified model of mammalian BCL-2 protein family interactions at the mitochondria." *Molecular Cell* 44.4 (Nov. 2011), pp. 517–531.
- Lopez, C. F., J. L. Muhlich, J. A. Bachman, and P. K. Sorger (2013). "Programming biological models in Python using PySB." *Molecular Systems Biology* 9, p. 646.
- Lovell, J. F., L. P. Billen, S. Bindner, A. Shamas-Din, C. Fradin, B. Leber, and D. W. Andrews (2008). "Membrane binding by tBid initiates an ordered series of events culminating in membrane permeabilization by Bax". *Cell* 135.6 (Dec. 2008), pp. 1074–1084.

- Maas, C., I. Verbrugge, E. d. Vries, G. Savich, L. W. v. d. Kooij, S. W. G. Tait, and J. Borst (2013). "Smac/DIABLO release from mitochondria and XIAP inhibition are essential to limit clonogenicity of Type I tumor cells after TRAIL receptor stimulation". *Cell Death and Differentiation* 17.10 (Mar. 2013), pp. 1613–1623.
- Maione, P., C. Gridelli, T. Troiani, and F. Ciardiello (2006). "Combining targeted therapies and drugs with multiple targets in the treatment of NSCLC." *The oncologist* 11.3 (Mar. 2006), pp. 274–284.
- Mallavarapu, A., M. Thomson, B. Ullian, and J. Gunawardena (2009). "Programming with models: modularity and abstraction provide powerful capabilities for systems biology". *Journal of The Royal Society Interface* 6.32 (Mar. 2009), pp. 257–270.
- Martinez-Caballero, S., L. M. Dejean, M. S. Kinnally, K. J. Oh, C. A. Mannella, and K. W. Kinnally (2009). "Assembly of the Mitochondrial Apoptosis-induced Channel, MAC". *The Journal of biological chemistry* 284.18 (May 2009), pp. 12235–12245.
- "Mechanism matters." (2010). *Nature medicine* 16.4 (Apr. 2010), p. 347.
- Moldoveanu, T., C. R. Grace, F. Llambi, A. Nourse, P. Fitzgerald, K. Gehring, R. W. Kriwacki, and D. R. Green (2013). "BID-induced structural changes in BAK promote apoptosis." *Nature Structural & Molecular Biology* 20.5 (May 2013), pp. 589–597.
- Muzzey, D., C. A. Gómez-Urbe, J. T. Mettetal, and A. van Oudenaarden (2009). "A systems-level analysis of perfect adaptation in yeast osmoregulation." *Cell* 138.1 (July 2009), pp. 160–171.
- Prahalad, A., C. Sun, S. Huang, F. Di Nicolantonio, R. Salazar, D. Zecchin, R. L. Beijersbergen, A. Bardelli, and R. Bernards (2012). "Unresponsiveness of colon cancer to BRAF(V600E) inhibition through feedback activation of EGFR." *Nature* 483.7387 (Mar. 2012), pp. 100–103.
- Rehm, M., H. Dussmann, R. U. Janicke, J. M. Tavaré, D. Kogel, and J. H. M. Prehn (2002). "Single-cell fluorescence resonance energy transfer analysis demonstrates that caspase activation during apoptosis is a rapid process. Role of caspase-3". *The Journal of biological chemistry* 277.27 (July 2002), pp. 24506–24514.
- Roth, A. D. et al. (2010). "Prognostic role of KRAS and BRAF in stage II and III resected colon cancer: results of the translational study on the PETACC-3, EORTC 40993, SAKK 60-00 trial." *Journal of Clinical Oncology* 28.3 (Jan. 2010), pp. 466–474.
- Saito, M., S. J. Korsmeyer, and P. H. Schlesinger (2000). "BAX-dependent transport of cytochrome c reconstituted in pure liposomes." *Nature Cell Biology* 2.8 (Aug. 2000), pp. 553–555.
- Schaber, J., M. Flöttmann, J. Li, C.-F. Tiger, S. Hohmann, and E. Klipp (2011). "Automated ensemble modeling with modelMaGe: analyzing feedback mechanisms in the Sho1 branch of the HOG pathway." *PLoS ONE* 6.3, e14791.
- Schafer, B., J. Quispe, V. Choudhary, J. E. Chipuk, T. G. Ajero, H. Du, R. Schneiter, and T. Kuwana (2009). "Mitochondrial outer membrane proteins assist Bid in Bax-mediated lipidic pore formation." *Molecular Biology of the Cell* 20.8 (Apr. 2009), pp. 2276–2285.

- Schlesinger, P. H. and M. Saito (2006). "The Bax pore in liposomes, Biophysics". *Cell Death and Differentiation* 13.8 (Aug. 2006), pp. 1403–1408.
- Schwille, P. and S. Diez (2009). "Synthetic biology of minimal systems". *Critical Reviews in Biochemistry and Molecular Biology* (July 2009).
- Shamas-Din, A., S. Bindner, X. Chi, B. Leber, D. W. Andrews, and C. Fradin (2015). "Distinct lipid effects on tBid and Bim activation of membrane permeabilization by pro-apoptotic Bax." *The Biochemical journal* 467.3 (May 2015), pp. 495–505.
- Shamas-Din, A., S. Bindner, W. Zhu, Y. Zaltsman, C. Campbell, A. Gross, B. Leber, D. W. Andrews, and C. Fradin (2013). "tBid undergoes multiple conformational changes at the membrane required for Bax activation." *Journal of Biological Chemistry* 288.30 (July 2013), pp. 22111–22127.
- Shrager, J. (2003). "The fiction of function." *Bioinformatics (Oxford, England)* 19.15 (Oct. 2003), pp. 1934–1936.
- Sorger, P. K. (2005). "A reductionist's systems biology: opinion." *Current Opinion in Cell Biology* 17.1 (Feb. 2005), pp. 9–11.
- Subburaj, Y., K. Cosentino, M. Axmann, E. Pedrueza-Villalmanzo, E. Hermann, S. Bleicken, J. Spatz, and A. J. García-Sáez (2015). "Bax monomers form dimer units in the membrane that further self-assemble into multiple oligomeric species". *Nature Communications* 6 (Aug. 2015).
- Suzuki, M., R. J. Youle, and N. Tjandra (2000). "Structure of Bax: Coregulation of Dimer Formation and Intracellular Localization". *Cell*.
- Tait, S. W. G. and D. R. Green (2010). "Mitochondria and cell death: outer membrane permeabilization and beyond". *Nature Reviews Molecular Cell Biology* 11.9 (Sept. 2010), pp. 621–632.
- Tan, C., P. J. Dlugosz, J. Peng, Z. Zhang, S. M. Lapolla, S. M. Plafker, D. W. Andrews, and J. Lin (2006). "Auto-activation of the apoptosis protein Bax increases mitochondrial membrane permeability and is inhibited by Bcl-2". *The Journal of biological chemistry* 281.21 (May 2006), pp. 14764–14775.
- Thiele, I. et al. (2013). "A community-driven global reconstruction of human metabolism." *Nature Biotechnology* 31.5 (May 2013), pp. 419–425.
- Tyson, J. J. (2007). "Bringing cartoons to life." *Nature* 445.7130 (Feb. 2007), p. 823.
- Venter, J. C. et al. (2001). "The Sequence of the Human Genome". *Science (New York, NY)* 291.5507 (Feb. 2001), pp. 1304–1351.
- Vidović, D., A. Koleti, and S. C. Schürer (2014). "Large-scale integration of small molecule-induced genome-wide transcriptional responses, Kinome-wide binding affinities and cell-growth inhibition profiles reveal global trends characterizing systems-level drug action." *Frontiers in genetics* 5, p. 342.
- Vo, T.-T., J. Ryan, R. Carrasco, D. Neuberg, D. J. Rossi, R. M. Stone, D. J. Deangelo, M. G. Frattini, and A. Letai (2012). "Relative mitochondrial priming of myeloblasts and normal HSCs determines chemotherapeutic success in AML." *Cell* 151.2 (Oct. 2012), pp. 344–355.

- Volkman, N., F. M. Marassi, D. D. Newmeyer, and D. Hanein (2013). "The rheostat in the membrane: BCL-2 family proteins and apoptosis." *Cell Death and Differentiation* (Oct. 2013).
- Weiss, J. N. (1997). "The Hill equation revisited: uses and misuses." *The FASEB Journal* 11.11 (Sept. 1997), pp. 835–841.
- Willis, S. N., L. Chen, G. Dewson, A. Wei, E. Naik, J. I. Fletcher, J. M. Adams, and D. C. S. Huang (2005). "Proapoptotic Bak is sequestered by Mcl-1 and Bcl-xL, but not Bcl-2, until displaced by BH3-only proteins". *Genes & Development* 19.11 (June 2005), pp. 1294–1305.
- Willis, S. N. et al. (2007). "Apoptosis initiated when BH3 ligands engage multiple Bcl-2 homologs, not Bax or Bak". *Science (New York, NY)* 315.5813 (Feb. 2007), pp. 856–859.
- Wingreen, N. and D. Botstein (2006). "Back to the future: education for systems-level biologists". *Nature Reviews Molecular Cell Biology* 7.11 (Nov. 2006), pp. 829–832.
- Yan, M., B. A. Parker, R. Schwab, and R. Kurzrock (2014). "HER2 aberrations in cancer: implications for therapy." *Cancer treatment reviews* 40.6 (July 2014), pp. 770–780.
- Youle, R. J. and A. Strasser (2008). "The BCL-2 protein family: opposing activities that mediate cell death". *Nature Reviews Molecular Cell Biology* 9.1, pp. 47–59.

Let us change our traditional attitude to the construction of programs: Instead of imagining that our main task is to instruct a computer what to do, let us concentrate rather on explaining to human beings what we want a computer to do.

Donald Knuth (Knuth, 1984)

2

Programming biological models in Python using PySB

2.1	Introduction	25
2.2	Results	31
2.3	Discussion	48
2.4	Methods	55
2.5	Supplementary information	66

AUTHORSHIP CONTRIBUTION

For this work I contributed to the software development of PySB (the project having been initiated by J. Muhlich and C. Lopez). I also developed the set of models in the EARM 2.0 repository and formulated several use cases for PySB, including the various types of model reuse. Specific software contributions to PySB are detailed on GitHub at <https://github.com/pysb/pysb/graphs/contributors>, and to EARM 2.0 at <https://github.com/sorgerlab/earm/graphs/contributors>. I took the primary role in preparing the figures and writing the manuscript.

ABSTRACT

Mathematical equations are fundamental to modeling biological networks but as networks get large and revisions frequent it becomes difficult to manage equations directly or to combine previously developed models. Multiple simultaneous efforts to create graphical standards, rule-based languages and integrated software workbenches aim to simplify biological modeling but none fully meets the need for transparent, extensible, and reusable models. In this paper we describe PySB, an approach in which models are not only created using programs, they are programs. PySB draws on programmatic modeling concepts from *little b* and ProMot, the rule-based languages BNGL and Kappa and the growing library of Python numerical tools. Central to PySB is a library of macros encoding familiar biochemical actions such as binding, catalysis and polymerization making it possible to use a high-level, action-oriented vocabulary to construct detailed models. As Python programs, PySB models leverage tools and practices from the open-source software community, substantially advancing our ability to distribute and manage the work of testing biochemical hypotheses. We illustrate these ideas using new and previously published models of apoptosis.

2.1 INTRODUCTION

Mechanistic studies that couple experimentation and computation typically rely on models optimized for specific questions and biological settings. Such “fit-to-purpose” models can effectively describe and elucidate complex biological processes but given available data they are usually restricted in scope, encom-

passing only a subset of cellular biochemistry (Schleich et al., 2012; Xu et al., 2010; Batchelor et al., 2008; Huber et al., 2011). Even in disciplines in which modeling is more mature, all-encompassing models are rare. However, it is common for fit-to-purpose models to require modification involving the addition, change or elimination of species and reactions based on new discoveries. Often a family of models is needed (Albeck et al., 2008b; Rehm et al., 2009; Xu et al., 2010; Muzzey et al., 2009; Kleiman et al., 2011) each of which represents a competing molecular hypothesis derived from the literature, a different way of encoding ambiguous “word models” drawn from molecular biology, or postulated differences in a network from one cell type to the next (Gnad et al., 2012). One of the promises of systems biology is that collaborative and iterative construction and testing of models can improve hypotheses by subjecting them to an ongoing process of testing and improvement. However, the current proliferation of independently derived fit-to-purpose models frustrates this goal (Krakauer et al., 2011). We require “second generation” tools that leverage existing approaches to biological model construction and documentation while adding new means for modifying, comparing and sharing models in a transparent manner.


Dynamic biological systems are generally modeled using stochastic or deterministic rate equations. The latter can be described using networks of ordinary differential equations that precisely represent mass action kinetics. However, when a network model has many species and variables, equations become a poor tool for model development and understanding. Even familiar biochemical processes are remarkably complex at the level of equations: for example, fully accounting for the binding and post-translational modifications underlying activation of growth factor receptors can require thousands of equations that are tedious to generate, hard to error-check and difficult to understand (Hlavacek et al., 2006). The need for frequent updates is also a challenge because even a simple modification of a biochemical cascade can require dozens of small changes in the corresponding equations. When operating at the level of equations it is also difficult to reuse the work of others directly. For example, in the field of receptor-mediated apoptosis, Howells et al. (2010) described a conceptual extension of a previously published model of Bcl-2 proteins (which control mitochondrial outer membrane permeabilization or MOMP, Box 2.1) (Chen et al., 2007a) by adding the BH₃-only Bcl-2 family member Bad. Interactions among the core set of Bcl-2 proteins were identical between the two models, but Howells et al. rewrote the original set of ODEs

simply to add a few new reactions. Manually re-building earlier models is not only time-consuming but also error-prone: as described in detail below, the practice has introduced errors and unintended changes in another pair of related apoptosis models. Moreover, the tendency to make numerous trivial changes in duplicated elements (e.g. by renaming species) makes it difficult to focus on key differences, frustrating later attempts at model comparison (Mallavarapu et al., 2009).

Several projects are underway to address the problems of model proliferation and complexity using formalisms that aim for abstraction, transparency and reusability. Chief among these is rule-based modeling (Hlavacek et al., 2006; Bachman and Sorger, 2011) in which models are created using specialized languages such as BioNetGen Language (BNGL) (Faeder et al., 2009) or Kappa (Danos et al., 2007b). These languages describe local interaction rules between specific domains or “sites” on molecules (e.g., enzymes and their substrates) and are easier to understand and reuse than equations (Danos, 2009). Rule-based approaches enable modeling of otherwise intractably complex systems in which post-translational modification and the formation of multi-protein signaling complexes give rise to large numbers of distinct species (Sneddon et al., 2011; Deeds et al., 2012; Blinov et al., 2006). Rules can also be subjected to formal analysis (Danos et al., 2008; Feret et al., 2009) and used to generate both deterministic and agent-based simulations (Danos et al., 2007a; Sneddon et al., 2011).

While powerful, rule-based languages such as BNGL and Kappa do not exploit “higher-order” patterns in biochemical systems such as multi-step catalysis, scaffold assembly, polymerization, receptor internalization, etc. These patterns often recur several times in a single model and also are frequently encountered in models of different molecular networks. Currently it is necessary to regenerate the rule sets for biochemical patterns each time the patterns arise. The creation of models having related but variant topologies presents an important special case of a higher-order pattern, particularly when a core process remains the same across all models and modifications focus on specific reactions. In both rule- and ODE-based models it is necessary to implement the change for each model in the set, a laborious process when the number of models is large. Modeling tools that leverage approaches from software engineering are one way to increase reusability and reduce duplication (Mallavarapu et al., 2009; Gnad et al., 2012; Pedersen and Plotkin, 2008; Danos et al., 2009; Mirschel et al., 2009). In particular, LISP frameworks such

Box 2.1: TRAIL-mediated apoptosis and the Bcl-2 protein family.



TRAIL is a prototypical pro-death ligand that binds transmembrane DR4 and DR5 receptors and leads to formation of intracellular, multi-component death inducing signaling complexes (DISCs). Autocatalytic processing of initiator procaspases-8 and 10 at the DISC allows the enzymes to cleave procaspase-3 but caspase activity is held in check by XIAP, an E3 ubiquitin ligase that blocks the caspase-3 active site and targets the enzyme for ubiquitin-mediated degradation. In most cell types, activation of caspase-3 and consequent cell killing requires mitochondrial outer membrane permeabilization (MOMP). MOMP allows translocation of cytochrome *c* and Smac into the cytosol where Smac binds and inactivates XIAP and cytochrome *c* activates pro-caspase-9, two reactions that generate cleaved, active caspase-3.

MOMP is regulated by a family of ~20 Bcl-2 proteins (Youle and Strasser, 2008) having three functions: pro-apoptotic effectors Bax and Bak assemble into pores, anti-apoptotic Bcl-2, Bcl-XL, Mcl-1, Bcl-W and A1 proteins block pore formation and the “BH3-only proteins” such as Bid, Bim, and Puma activate the effectors and inhibit the anti-apoptotics. Bid, the most important BH3-only protein for extrinsic cell death, is a direct substrate of caspases 8/10 and its active form (tBid) binds to and activates Bax and Bak via a recently elucidated structural transition (Gavathiotis et al., 2010; Kim et al., 2009). The overall pathway can be roughly divided into a “receptor to Bid” module (yellow in the figure), a “pore to PARP” module (blue), and a MOMP module (orange).

Structural and cellular studies of Bcl-2 proteins are consistent with a variety of different mechanisms for MOMP. A direct model posits that effectors form pores only when activated by proteins such as tBid (Letai et al., 2002; Kim et al., 2006; Ren et al., 2010). The indirect model proposes that Bax and Bak are constitutively able to form pores but are held in check by anti-apoptotic Bcl2 proteins (Willis et al., 2007). Recent studies support a combination of both direct and indirect mechanisms (Mérino et al., 2009; Llambi et al., 2011; Leber et al., 2010). The “embedded together” model emphasizes the active role of membranes in determining the conformational states and activity of Bcl-2 proteins and that the anti-apoptotic Bcl2 proteins possess all of the same functional interactions as the effectors except pore formation and therefore function as dominant-negatives (Billen et al., 2008; Leber et al., 2010). Controversy about MOMP mechanisms reflects the large number of Bcl-2 proteins involved, the role of protein compartmentalization and localization in activity, the diversity of apoptosis inducers and the fact that different cell types express different Bcl-2 proteins at very different levels. Detailed mechanistic models of MOMP are nonetheless important for rationalizing the selectivity of anti-Bcl2/Bcl-XL drugs such as ABT-263, understanding the oncogenic potential of proteins such as Mcl-1 and Bcl-2, and elucidating the precise mechanisms of action of pro-apoptotic chemotherapeutics.

as *little b* (Mallavarapu et al., 2009) and ProMot (Mirschel et al., 2009) have demonstrated the value of programmatic approaches. However, ProMot does not use rules, limiting its effectiveness for combinatorially complex systems; *little b*, while implementing rules internally, does not interoperate with tools and languages from the broader rule-based modeling community and is no longer in development (the similarities and differences between the *little b* and ProMot approaches have been described previously (Mallavarapu et al., 2009)). Combining the strengths of rule-based and programmatic approaches to modeling is a key goal of the work described here.

A benefit of modeling biological systems using contemporary approaches from computer science and open-source software engineering is the ready availability of tools and best practices for managing and testing complex code. Good software engineering practice promotes abstraction, composition and modularity (Mallavarapu et al., 2009; Mirschel et al., 2009). Through abstraction, the core features of a concept or process are separated from the particulars: for example, a pattern of biochemical reactions (e.g., phosphorylation-dephosphorylation of a substrate) is described once in a generic form as a subroutine and then instantiated for specific models simply by specifying the arguments (e.g., species such as Raf, PP2A, and MEK). In programming, abstraction is achieved through the use of parameterizable functions or macros that are written once and then invoked as needed. Functions can be built up from other functions, a process known as composition. Abstraction and composition can occur at all levels of complexity: just as complex functions can be built from simple functions, large programs can be built up from smaller subsystems that are documented and tested individually. When these subsystems have well-defined input-output interfaces, they can be used as libraries that can make it possible to write new programs using a simple vocabulary of well-tested concepts (e.g., a library of biochemical actions or core pathways such as the MAPK cascade) (Pedersen and Plotkin, 2008). The decomposition of complex biological models in this fashion facilitates extensibility and transparency, since well-developed mechanisms can be reused and changes can be localized to the subsystem that needs revision.

Contemporary software engineering has much to teach us about the difficult task of developing and documenting models in a distributed setting. Software engineers “publish” their findings using robust programming tools that support code annotation, documentation, and verification, all significant chal-

lenges in biological modeling (Hlavacek, 2009). The open source software community also provides a valuable socio-cultural framework for managing large, collaborative projects in the public domain. Version control tools such as Git and Subversion, along with “social coding” websites such as GitHub, have facilitated the collaborative development of software as complex as the kernel of the Linux operating system (<http://github.com>). It would be highly desirable to exploit such social and technical innovation in solving the problems of incremental model development and reuse in biology.

In this paper we describe PySB, an open-source programming framework written in Python that allows concepts and methodologies from contemporary software engineering to be applied to the construction of transparent, extensible, and reusable biological models (<http://python.org>; Oliphant, 2007). A critical feature of modeling with PySB is that models are Python programs, and tools for documentation, testing, and version control (e.g, Git) can be used to help manage model development. Strictly speaking, a PySB “model” is a Python program, that, when executed, produces another program (the underlying reaction rules) that can be analyzed or used to create equations for simulation. The PySB framework provides a high-level action-oriented vocabulary congruent with our intuitive understanding of biochemistry (A phosphorylates B, C translocates to the nucleus, etc.). PySB is closely integrated with Python numerical tools for simulation and parameter estimation and graphical tools that enable plotting of model trajectories and topologies. We demonstrate the use of PySB to re-instantiate 15 previously published ordinary differential equation (ODE)-based models of the Bcl-2 family proteins that regulate apoptosis (Albeck et al., 2008b; C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008; Howells et al., 2010). We show how PySB can be used to decompose models into reusable macros that can be independently tested and we generate composite models that combine a prior model from our lab describing extrinsic apoptosis (Albeck et al., 2008b; Albeck et al., 2008a) with alternative hypotheses about Bcl-2 family interactions from other investigators. Finally, we develop and calibrate an expanded cell death model that spans the diversity of the multi-protein Bcl-2 family and draws on findings from leading biochemists in the field, depicting a unified, “embedded together” mechanism for mitochondrial membrane permeabilization (Leber et al., 2010; Llambi et al., 2011). All models in this paper, along with the PySB source code and user manual, are available for sharing and further development at GitHub and

Box 2.2: PySB embeds a biological modeling language within Python.

Existing computer languages developed for biological modeling (e.g., BNGL or SBML) use a specialized syntax to concisely encode the detailed specifications of biological models. However, such languages, more generally known as domain-specific languages (DSLs), lack many features found in general-purpose programming languages (functions, classes, loops, etc.) that can be used for organizing complex code and making it more human-readable. They also often lack ancillary tools commonly found in general-purpose programming languages, such as testing frameworks or documentation generation support. Models written using PySB are programs in the Python programming language—they are not specialized file formats interpreted by modeling programs. When executed, a PySB model programmatically “builds up” the elements of a rule-based model (molecule types, rules, parameters, etc.) until the specification is complete; the Python model object that results can then be subject to further analysis. A traditional, object-oriented approach to building up a model in this way would require many extra lines of code to create and track objects relative to a DSL. PySB streamlines the process of programmatically building models by overloading several Python operators (+, >>, <>, (), %) to allow biological rules to be expressed in a chemistry-like syntax based on BNGL. In addition, the SelfExporter helper class allows models to be built up declaratively in BNGL-like fashion, further minimizing the required code. The overall result is a specialized language for biological modeling “contained within” Python and implemented using Python operators (the Python package SymPy for symbolic mathematics uses a similar approach). While this “syntactic sugar” streamlines the most common modeling use cases, in some advanced modeling scenarios a more traditional programming syntax may be preferred; the streamlined syntax is thus entirely optional. The PySB syntax is described in detail in the [Supplementary information](#).

the PySB website (<http://pysb.org>; [Methods](#)).

2.2 RESULTS

We chose Python as the language for PySB due to its widespread use in the computational biology community, support for object-oriented and functional programming, and rich ecosystem of mathematical and scientific libraries. At the outset we determined that PySB should interoperate seamlessly with BioNetGen ([Faeder et al., 2009](#)) and Kappa ([Danos et al., 2007b](#)) and thereby build on substantial investments in rule-based modeling. A PySB model consists of instances of a core set of classes: `Model`, `Monomer`, `Parameter`, `Compartment`, `Rule`, `Initial`, and `Observable`, closely mirroring the form of BNGL and Kappa models. However, in PySB the component declarations return software objects inside Python, allowing model elements to be manipulated programmatically. To simplify the process of writing rules and to maximize the syntactic match with BNGL, PySB redefines (overloads) some of the mathematical operators in Python to create a shorthand that resembles chemistry notation ([Box 2.2](#)). For example, in the context of a PySB rule definition, the “+” operator (which in other contexts represents mathematical addition) is used to enumerate a list of reactants or products. It is not necessary to use overloaded PySB operators but it makes the models easier to write and understand (see “PySB Syntax” in [Methods](#) and the [Supplementary information](#)).

By way of illustration, consider a “*Hello World*” program in PySB involving reversible binding of proteins L and R, each of which contains a single binding site *s*. A PySB program for this simple reaction has high overhead relative to the equivalent set of ODEs but serves to introduce the basic PySB syntax and show how it interoperates with other software such as BNG, the VODE integrator (Brown et al., 1989), and the Matplotlib plotting library (Hunter, 2007). In the first block of the program (Figure 2.1A; block 1) a declaration of molecule types (“monomers” in PySB) is followed by the forward and reverse rate parameters, *k_f* and *k_r*, and initial conditions for unbound L and R. The syntax for the reversible binding rule (which will be familiar to users of rule-based languages) reads as follows: when the proteins L and R both have empty binding sites *s* (e.g. `L (s=None)`), they reversibly bind to form a complex that shares a single “bond” (e.g., `L (s=1) % R (s=1)`), at rates *k_f* and *k_r*. This approach to naming binding sites (and calling interactions “bonds,” without implying covalent modification) is drawn from rule-based languages and makes it possible to describe molecules having multiple sites of interaction with different specificities, modifications, and occupancy states (e.g. the distinct binding sites on the TRAIL receptor for ligand and DISC adaptor proteins). The “*Hello World*” model concludes by designating an observable of interest, the complex LR. (For additional details on the syntax used in this model, see Box 2.2 and the “PySB Syntax” section of the Supplementary information; a UML class diagram of the core PySB classes is also provided in Supplementary Figure S3.) The second block of code in Figure 2.1A defines a time range and calls the PySB function `odesolve`, which performs deterministic model simulation by invoking BNG (to generate the reaction network) and the numerical integrator VODE (see also Figure 2.1B). Simulation results are returned as a matrix which is graphed using the `plot` command from Matplotlib (Figure 2.1A, block 3; Figure 2.1B).

2.2.1 USING MACROS TO MODEL RECURRENT BIOCHEMICAL ACTIONS

The benefits of PySB become apparent only with more complex and interesting models in which programming constructs such as conditionals, loops, functions, classes, and modules are used to define reusable elements. The complexity of these elements can vary from a few species to multi-component cascades. *Macros*, reusable Python functions that are programmatically expanded into rules and reactions, define




Figure 2.1: Creation and simulation of a “Hello World” model in PySB.

(A) Creation and deterministic simulation of a model using PySB. The call to `Model()` creates the `pysb.core.Model` object to which all subsequently declared components are added. The first block of code declares the molecule types, parameters, initial conditions, reversible reaction rule and observable necessary for modeling and simulating the reversible binding of proteins L and R. The second block of code calls the `odesolve` function from the `pysb.integrate` module to generate and integrate the ODEs. The third block plots the simulated timecourse using the Matplotlib `plot` command. Numbers associated with the code blocks identify the correspondences between the code and the control flow shown in (B).

(B) Control flow for an ODE simulation of a PySB model. The columns “User,” “PySB,” and “External Tools” indicate the locus of control of each step in the process (boxes). The “Result” column shows the result of each individual step: the gray box indicates results of steps that are internal to the call to `odesolve`, while the other results are visible to the user at the top level. After declaring the model elements as in (A), the user calls `odesolve`, which generates the corresponding BNGL for the model, invokes BNG externally on the generated BNGL code to create the reaction network, parses the reaction network to generate the corresponding set of ODEs, and calls an external integrator (VODE) to generate the trajectories. The trajectories are returned to the user as a Numpy record array, where they are visualized with a call to the Matplotlib function `plot`.

low-level biochemical actions such as “catalyze,” “bind,” or “assemble,” mirroring the way we describe biological processes verbally. PySB currently contains 13 general-purpose macros covering reversible binding, catalytic modification, synthesis, degradation, and pore assembly. Users can generate other model-specific macros to implement new or uncommon mechanisms, thereby creating libraries of biochemical actions for subsequent modeling projects (we are currently adding new macros ourselves); implementation of distributed or coordinated phosphorylation/dephosphorylation cascades and loops is an obvious candidate for such a library.

As an example, the `catalyze` macro implements a mass action model of an enzyme-mediated biochemical transformation (Figure 2.2A) based on six user-specified arguments: the enzyme and its binding site for substrate, the substrate and its binding site for enzyme, the product, and a list of forward, reverse,

and catalytic rate constants. [Figure 2.2A](#) illustrates the application of `catalyze` to a reaction in which caspase-8 cleaves Bid to form truncated Bid (tBid; see [Box 2.1](#) for a description of the relevant biology). Improved transparency is an important benefit of using macros in that they explicitly describe how elementary reactions are implemented. For example, `catalyze` invokes a two-step model of catalysis in which an enzyme-substrate complex is formed as an intermediate step: $E + S \rightleftharpoons ES \rightarrow E + P$. Some published models of apoptosis ([Cui et al., 2008](#); [C. Chen et al., 2007a](#); [C. Chen et al., 2007b](#); [Howells et al., 2010](#)) use a pseudo-first-order, one-step approximation $E + S \rightarrow E + P$ that has the merit of fewer parameters. However, depending on time scales and parameter values there is a profound difference in the dynamics of one- and two-step catalysis: the one-step model, for example, makes the strict assumption that the enzyme always operates in its linear range and cannot be saturated. These differences are not apparent from the visual or existing verbal descriptions of the ODEs but instead require careful retrospective analysis. In contrast, PySB allows exploration of mechanistic differences simply by calling an alternative catalysis macro, `catalyze_one_step` ([Supplementary Figure S1B](#)). The benefit of macros is not only that the model is more concise (the `catalyze` macro call replaces two rules and four ODEs) but also that the difference between one- and two-step catalytic schemes is clearly declared and need not be inferred retrospectively. Since macros are tested programmatically they also ensure correct instantiation of forward, reverse and catalytic reactions, an important benefit since failure to implement these processes correctly is remarkably common (see below and [Supplementary note](#)).

The power of macros is most evident with complex biochemical actions. For example, the `assemble_pore_sequential` macro ([Figure 2.2B](#)) implements sequential assembly of a membrane-associated pore from multiple identical subunits. It was written to explore how Bax and Bak associate in the mitochondrial outer membrane to form the pore that translocates cytochrome *c* and Smac into the cytosol during MOMP ([Youle and Strasser, 2008](#)). The arguments to `assemble_pore_sequential` are the identity of the pore-forming protein and its two binding sites, the maximum number of proteins in a pore and a table containing rate constants for each assembly step. *In vitro* experiments suggest that Bax and Bak assemble by sequential addition of subunits ([Martinez-Caballero et al., 2009](#)) but the precise size of an active MOMP pore is unknown. Exploring the effects of changing the number *n* of subunits in a

A

Example Macro Call

Macro

123456

catalyze(C8, 'bf', Bid(state='U'), 'bf', Bid(state='T'), klist)

1: Enzyme (Caspase-8)4: Substrate binding site name

2: Enzyme binding site name5: Product (T for truncated, cytosolic Bid)

3: Substrate (full length; cytosolic Bid)6: List of rate parameters

Implementation

Supplementary Figure S1A

BNGL Rules

C8(bf) + Bid(bf,state~U) <-> C8(bf!1).Bid(bf!1,state~U) kf, kr
C8(bf!1).Bid(bf!1,state~U) -> C8(bf) + Bid(bf,state~T) kc

ODEs

C8: ds0/dt = kc*s2 - kf*s0*s1 + kr*s2
Bid: ds1/dt = -kf*s0*s1 + kr*s2
C8_Bid: ds2/dt = -kc*s2 + kf*s0*s1 - kr*s2
tBid: ds3/dt = kc*s2

B

Example Macro Call

Macro

12345

assemble_pore_sequential(Bax, 's1', 's2', 6, ktable)

1: Subunit name4: Maximum number of subunits in pore

2, 3: Sites for binding neighboring subunits5: Table of rate parameters

BNGL Rules

5 rules

ODEs

6 ODEs

C

Example Macro Call

Macro

Arguments

1

bind_table([[Bcl2, BclXL, BclW, Mcl1, Bfl1],
[Bid, 66, 12, 10, 10, 53],
[Bim, 10, 10, 38, 10, 73],
[Bad, 11, 10, 60, None, None],
[Bik, 151, 10, 17, 109, None],
[Noxa, None, None, None, 19, None],
[Hrk, None, 92, None, None, None],
[Puma, 18, 10, 25, 10, 59],
[Bmf, 24, 10, 11, 23, None]],
'bf', 'bf', kf=1e-3])

234

1: Bind table data (molecules and rates)3: Binding site for column-group proteins

2: Binding site for row-group proteins4: Default association rate

BNGL Rules

28 rules

ODEs

41 ODEs

Figure 2.2: Three examples of mechanistic abstractions using macros. Full implementation of all macros can be found in the `macros.py` file in the PySB source code online (Methods).

(A) `catalyze`. The example call shows how the `catalyze` macro is called to add a catalytic reaction in which active caspase-8 (C8) binds to untruncated Bid (`Bid(state='U')`) to yield tBid (`Bid(state='T')`). Rate parameters (forward, reverse, and catalytic) are provided in the list `klist`. The “Basic implementation” in [Supplementary Figure S1A](#) shows the Python source code for a simplified version of the `catalyze` macro. Execution of the `catalyze` macro leads to the creation and addition of two rules to the model, which, when converted into BNGL, take the form shown below. Generation of the reaction network via BNG then results in the system of four ODEs shown at bottom.

(B) `assemble_pore_sequential` models the assembly of a pore in a sequential fashion in which monomers bind to form dimers, dimers bind monomers to form trimers, trimers bind monomers to form tetramers, etc. Pores of size three (trimers) and above have a closed, ring-shaped topology, reflecting the variable structure and stoichiometry for the Bax pore suggested by recent *in vitro* experiments ([Martinez-Caballero et al., 2009](#)). The maximal size of the pore is given by the fourth argument to the macro (i.e., 6). With the parameters shown in the example, the execution of the macro results in five rules (for binding of monomers to monomers, monomers to dimers, monomers to trimers, etc.) and six species and ODEs (monomers through hexamers). Pores with greater stoichiometry can be modeled simply by changing the pore size argument in the macro call.

(C) `bind_table` is used to concisely represent combinatorial binding among two related groups of molecules. In the example call the species in the column headers are the five known anti-apoptotic Bcl-2 proteins, while the row headers are various pro-apoptotic BH3-only proteins. The table entries represent the dissociation constants for binding between each pair of proteins drawn from *in vitro* measurements, given in units of nanomolar ([Certo et al., 2006](#); [Willis et al., 2005](#)); the Python constant `None` indicates that no binding occurs. (In place of a dissociation constant, a table cell may alternatively contain a pair of explicit association and dissociation rates.) The names of the binding sites for the row-group and column-group proteins (i.e., “bf”) are given as the second and third arguments. The final argument (i.e. `kf=1e-3`) specifies a default association rate to be applied to all reactions, given in units of nanomolar⁻¹ second⁻¹. The execution of the `bind_table` call results in the instantiation of 28 reversible binding rules, each with the given association rate and a dissociation rate calculated from the dissociation constant provided in the table entries; this further expands to 41 ODEs.

pore requires creation and modification of $n - 1$ BNGL rules or n ODEs but only one value in the PySB `assemble_pore_sequential` macro.

Many biochemical processes are controlled by families of proteins that have overlapping binding specificities. In apoptosis, MOMP is regulated by ~ 20 pro- and anti-apoptotic Bcl-2 family proteins that have a range of affinities for each other and bind in various arrangements. Modeling the binding of any two Bcl-2 proteins is simple, but managing equations or rules for all possible binding reactions is much harder. The `bind_table` macro uses a simple tabular representation to model interactions among members of multi-protein families (Figure 2.2C). The first argument to the macro is a table (a list of lists in Python) in which row and column headers identify pairs of interacting proteins and each table entry contains binding constants or the value `None` to indicate that there is no measurable interaction. The second and third arguments specify the binding site names for row and column species, respectively. In addition to being concise (a single `bind_table` call for the Bcl-2 proteins generates 28 rules and 41 ODEs) the tabular format highlights relationships between proteins by grouping them into functional classes; new family members can be introduced simply by adding rows and columns. A `bind_table` is therefore a simple computable representation of the “binding codes” that have been created by others to summarize structural and biochemistry studies on Bcl-2 family members (Certo et al., 2006; Kuwana et al., 2005; L. Chen et al., 2005). By changing the first argument in the `bind_table` call it is straightforward to explicitly substitute one set of binding data for another, a useful feature for exploring differences in published binding codes or for modeling the behavior of mutant proteins (Debartolo et al., 2012; Fire et al., 2010). Models that use macros such as `bind_table` naturally acquire a “self-documenting” character that minimizes the need for additional explanatory description (see Figures 2.4B and 2.5A for examples of the use of `bind_table` in the context of models of MOMP initiation).

2.2.2 MODULES AND MODEL REUSE

Macros abstract biochemical reactions at a fairly low level of detail involving a few proteins but Python also supports a wide variety of methods for reusing more complex model elements. We have found three to be particularly useful: (1) *Instance Reuse*, for making small changes to an existing model; (2) *Module*

Reuse, for programmatically composing a model from reusable pieces or modules; and (3) *Class Reuse*, for models involving combinatorial variation of several independent features. Of these, Instance Reuse is the simplest, entailing programmatic duplication of a previous model and explicit specification of new and modified elements. Instance Reuse proved to be an appropriate way to update the PySB version of EARM 1.0 (Albeck et al., 2008a) to include synthesis and degradation reactions (Figure 2.3A). This approach replaces conventional “cut-and-paste” copying and editing of ODEs or BNGL rules. Reuse is achieved by cloning the old model into a new model and then explicitly declaring the elements that are added or modified, making the changes easy to understand and track.

Module Reuse involves the separation of a model into independent elements (“motifs” or “modules”; see also Figure 2.7) that are written as callable subroutines in Python. It is not necessary for a biological process to be modular in a functional sense for modularization in PySB to be advantageous. Building a model from subroutines enables a “mix-and-match” approach in which a subset of the interactions in a model is subjected to revision or re-examination while the rest remain the same. For example, we divided EARM 1.0 into three modules each involving self-contained blocks of PySB code for: (1) reactions from ligand-death receptor association to binding of DISC components, (2) interactions among Bcl-2 family members controlling MOMP, and (3) the cascade of reactions involving initiator and effector caspases and their immediate regulators (Figure 2.3B; Box 2.1). A series of papers examining alternative models of MOMP have been published by multiple groups (C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008; Howells et al., 2010) but MOMP has most commonly been studied in isolation from reactions occurring upstream and downstream. However, it has been shown that multi-protein cascades do not exhibit the same behavior in isolation as when they are part of larger networks (W. W. Chen et al., 2009; Del Vecchio et al., 2008). One of the primary aims of modeling signal transduction is to contextualize molecular mechanisms by embedding them in a network context. Thus, studies of extrinsic apoptosis would benefit from models in which alternative hypotheses for MOMP regulation are embedded in a more complete reaction pathway. Using conventional modeling tools it is challenging to add MOMP “mini-models” to upstream and downstream reactions (Albeck et al., 2008b). In contrast, in PySB this type of composition is simple: we have written a PySB program in which any of 15 models of MOMP are




Figure 2.3: Three approaches to model re-use. Boxes with gray tabs represent Python modules/files. Statements marked “Reuse” identify the point of re-use of previously created model code. Background highlighting indicates correspondence between elements in the re-use schematic and the PySB code.

(A) Direct re-use and subsequent modification of pre-existing model code. A pre-existing model is declared in its own Python file (`earm_1_0.py`). The extending model, in the file `earm_1_3.py` (representing a later version) imports and duplicates the model object from `earm_1_0.py` and subsequently adds a list of synthesis and degradation reactions.

(B) Reuse of modules using macros. Macros instantiating the components for the upstream (`rec_to_bid`) and downstream (`pore_to_parp`) portions of the extrinsic apoptosis pathway are placed in a Python file, `earm_modules.py`. Variant models differing only in the reaction topology for MOMP initiation (e.g., `model_variant.py`) are then created by invoking these macros for the shared upstream and downstream elements.

(C) Reuse and recombination of model elements through class inheritance. A shared implementation of the compartmentalization-independent reactions for Bax activation (i.e., `tBid_activates_Bax`) is contained within the base class `CptBase`. Alternative compartmentalization strategies are implemented in the child classes `TwoCpt` and `MultiCpt`, which separately implement the compartmentalization-dependent reactions for tBid and Bax translocation (i.e., `translocate_tBid_Bax`). Because these classes inherit from `CptBase`, they acquire the implementation of `tBid_activates_Bax`, representing a point of reuse. Alternative protein interaction topologies are implemented within the `build_model` function in the two classes `Topo1` and `Topo2`. Models with either of the compartmentalization implementations (e.g., `TwoCpt`) and either of the interaction topologies (e.g., `Topo1`) can then be created dynamically by inheriting from the appropriate classes, representing an additional point of reuse. Readers familiar with the concept of polymorphism from object-oriented programming will note that calling the `build_model` method on any of the hybrid classes will polymorphically refer to the correct implementations in the parent classes.

called up along with common `rec_to_bid` and `pore_to_parp` modules to create 15 fully functioning hybrid models of extrinsic apoptosis (Figure 2.3B). The upstream and downstream reactions have also been modeled in several different ways by others (Schleich et al., 2012; Fricker et al., 2010; Neumann et al., 2010; Rehm et al., 2006; O'Connor et al., 2008) and it would be straightforward to use Module Reuse to combine different proposals for `rec_to_bid` and `pore_to_parp` with different MOMP modules.

Class Reuse is a third and more sophisticated approach that exploits the class inheritance mechanism in Python. Common model features are coded in a base class and model variants are written as “child” classes of the base class, able to inherit code from the base class directly with or without programmatic modification. Code from multiple variants can then be combined by further inheritance from more than one of these classes. For example, we used Class Reuse to model the effects of reaction compartmentalization on interactions among pro-apoptotic Bcl-2 proteins at the mitochondrial membrane (Figure 2.3C). In one case we assumed that reactions took place in two well-mixed reaction compartments corresponding to cytosol and membrane (Figure 2.3C, `TwoCpt`) and in the second we assumed that each mitochondrion constituted a distinct reaction compartment (`MultiCpt`). In addition, we independently explored different reaction topologies involving the Bcl-2 proteins `tBid` and `Bax`. Both topologies (`Topo1` and `Topo2`) include the translocation of `tBid` and `Bax` to membranes but only the second topology (`Topo2`) incorporates activation of `Bax` by `tBid`. We used inheritance to automate creation of four different models having different compartmentalization schemes and reaction topologies (e.g., `TwoCpt_Topo1`). The notable feature of Class Reuse is that model variants are created and combined over multiple independent “axes”—in this example compartmentalization and protein interaction topology—transparently and with no duplicated code.

2.2.3 INTEGRATION WITH THE PYTHON ECOSYSTEM AND EXTERNAL MODELING TOOLS

The iterative process of model development is dramatically accelerated when tools for model creation, simulation, analysis and visualization are integrated. Many commercial and academic software packages including Mathematica (Wolfram Research, Inc., n.d.) and MATLAB (Mathworks, n.d.) provide integrated tools for equation-based models but are unwieldy to use with rule-based or programmatic ap-

Tool	Reference	Interface	Description (relevance to PySB)
NumPy	(Oliphant, 2007)	Python	Efficient array and matrix operations
SciPy	(Oliphant, 2007)	Python	Scientific algorithms, e.g., ODE integration, statistics, and optimization
SymPy	(SymPy Development Team, 2012)	Python	Symbolic manipulation of mathematical expressions
Matplotlib	(Hunter, 2007)	Python	Plotting and other data visualizations
Graphviz	(Gansner and North, 2000)	Python	Layout and rendering of node-edge graphs
BNG	(Faeder et al., 2009)	Wrapper	Translation of rules to a reaction network; stochastic simulation
Kappa	(Danos et al., 2007b)	Wrapper	Stochastic simulation; visualization and analysis of rule models
SBML	(Hucka, 2003)	Export	Compatibility with SBML tools
Mathematica	(Wolfram Research, Inc., n.d.)	Export	General-purpose scientific computing
MATLAB	(Mathworks, n.d.)	Export	General-purpose scientific computing

Table 2.1: Integration with external modeling tools.

proaches because models must be exported and imported using SBML ([Hoops et al., 2006](#); [Maiwald and Timmer, 2008](#)). At the same time, rule-based model editors such as RuleBender for BNGL ([Smith et al., 2012](#)) and RuleStudio for Kappa (<https://github.com/kappamodeler/rulestudio>) facilitate development of rule-based models but do not incorporate tools for data analysis, parameter fitting, and symbolic math. Simply by virtue of being written in Python, PySB interacts natively with a large and growing library of open-source scientific software such as NumPy, SciPy, SymPy, and Matplotlib ([Table 2.1](#)). Models written using PySB can also exploit Python tools for documentation generation (`sphinx`) and for unit testing (`unittest`, `nose`, and `doctest`) both of which we used extensively in creating the models of extrinsic apoptosis described below.

To interface PySB with BNG and Kappa, which are not implemented in Python, we wrote Python “wrapper” libraries, providing access to agent-based simulation, static analysis and visualization. The wrappers also manage the syntactic differences between BNGL and Kappa, allowing either to be used for the same PySB model. Models can be trivially exported in BNGL format for use with established “all-in-one” tools that support BNGL, such as V-Cell ([Moraru et al., 2002](#)). PySB can export models as systems of ODEs in formats for SBML ([Hucka, 2003](#)), MATLAB, Mathematica, or PottersWheel ([Maiwald and Timmer, 2008](#)). Finally, to facilitate unique identification of model components, we added a lightweight annotation capability that allows any model element (including macros and modules) to be tagged with

identifiers from external databases using subject-object-predicate triples compatible with MIRIAM (Le Novère et al., 2005). The net result is a software environment that combines the flexibility of a general-purpose scientific computing package with programmatic and rule-based modeling tools and an open-source code base.

2.2.4 EARM VERSION 2.0: A FAMILY OF MODELS OF EXTRINSIC APOPTOSIS AND MOMP

To explore the ability of PySB to model the latest molecular data on apoptosis while also building on previous work, we used macros and Module Reuse to construct a family of cell death models involving 15 different schemes for MOMP regulation. Seven of the MOMP modules were previously published by the research group of Shen and colleagues (C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008), one module extends one of these models (Howells et al., 2010), five modules are from the work of Albeck et al. (Albeck et al., 2008b), and three modules are entirely new and incorporate more complete sets of interactions among Bcl-2 proteins (Figure 2.4A). The three new modules are derived from word models in recent studies from Green and Andrews that unify previously competing mechanisms of pore formation (Llambi et al., 2011; Billen et al., 2008; Leber et al., 2010). Each of the 15 modules was instantiated in PySB as a distinct subroutine that can be called and analyzed in the context of a receptor-to-caspase pathway. The set of 15 MOMP modules is by no means complete, and several noteworthy models of extrinsic apoptosis and MOMP (Legewie et al., 2006; Bagci et al., 2006; Bentele et al., 2004; Rehm et al., 2009; Dussman et al., 2009) have not yet been coded in PySB. However, since our objective is to explore model reuse and composition using PySB we limited ourselves to the 15 MOMP-focused examples described above. We collectively denote the resulting set of 30 variant models (15 models only of MOMP plus 15 models of extrinsic apoptosis incorporating the MOMP modules) as Extrinsic Apoptosis Reaction Model version 2.0 (EARM 2.0); the models are summarized in Table 2.2 and are available as a Python package with source code downloadable from GitHub (Methods).

While porting existing models into PySB we observed that several published ODE networks contained one or more errors relative to their verbal or graphical descriptions in the original paper (Cui et al., 2008) (see Supplementary note). We were unable to discern whether the errors in the published ODE networks




Figure 2.4: Refactoring of published models into PySB.

(A) Relationships between the models examined in this paper. The “Albeck Group” incorporates a series of incrementally expanded models shown in Figure 11 of Albeck et al. (2008b); the “Shen/Howells Group” incorporates models from three papers from the research group of Shen (C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008) and a derivative model from Howells et al. (2010); the “Lopez group” includes three expanded models introduced in this paper. The arrows indicate that one model has been derived or extended from a prior model and point in the direction Base Model → Derivative Model.

(B) The “Direct” model from Chen et al. (C. Chen et al., 2007b) in its original ODE-based representation.

(C) Conversion of the Chen “Direct” model to a PySB module. The execution of the `chen_febs_direct` function results in rules that exactly reproduce the ODEs shown above (the molecule type `Bad` in the PySB function corresponds to the generic enabler species `Ena` in the original equations; `Bid` corresponds to the generic activator `Act`). The macro `catalyze_one_step_reversible` implements the two-reaction scheme $E + S \rightarrow E + P, P \rightarrow S$; `assemble_pore_spontaneous` implements the order-4 reaction $4 \times subunit \leftrightarrow pore$. The `bind_table` macro is illustrated in Figure 2.2C.

(D) Model extension in PySB. Module Reuse (Figure 2.3B) was used to implement the “Direct” model from Cui et al. (2008) as an extension of the prior “Direct” model from Chen et al. shown in Figure 2.4C (C. Chen et al., 2007b). Invocation of the PySB function `chen_febs_direct` incorporates the elements of the original Chen et al. model; subsequent statements specify the modifications and additions required to yield the derived model from Cui et al. (2008).

Model name ^a	ID (full/MOMP)	Reference	MOMP-only (<i>Mnb</i>) ^b			Full apoptosis (<i>Mna</i>) ^c		
			Rules	ODEs	Parameters	Rules	ODEs	Parameters
Lopez Embedded	M1a/b	This paper	39	40	78	66	76	133
Lopez Direct	M2a/b	This paper	27	32	58	54	68	113
Lopez Indirect	M3a/b	This paper	29	34	64	56	70	119
Albeck 11b	M4a/b	(Albeck et al., 2008b)	7	13	17	34	48	71
Albeck 11c	M5a/b	(Albeck et al., 2008b)	11	17	25	38	52	79
Albeck 11d	M6a/b	(Albeck et al., 2008b)	12	18	27	39	53	81
Albeck 11e	M7a/b	(Albeck et al., 2008b)	14	21	31	41	56	85
Albeck 11f	M8a/b	(Albeck et al., 2008b)	14	21	31	41	56	85
Chen 2007 Biophys J	M9a/b	(C. Chen et al., 2007a)	6	7	12	37	49	75
Chen 2007 FEBS Direct	M10a/b	(C. Chen et al., 2007b)	5	8	12	36	48	74
Chen 2007 FEBS Indirect	M11a/b	(C. Chen et al., 2007b)	3	6	9	34	48	72
Cui Direct	M12a/b	(Cui et al., 2008)	18	10	26	49	52	89
Cui Direct 1	M13a/b	(Cui et al., 2008)	22	11	33	53	53	96
Cui Direct 2	M14a/b	(Cui et al., 2008)	23	11	34	54	53	97
Howells	M15a/b	(Howells et al., 2010)	14	12	22	45	49	84

^a Model names are drawn from the first author of the paper in which the mathematical model was published.

^b MOMP-only variants are identified as *Mnb*, e.g., M1b for the MOMP-only variant of “Lopez Embedded”.

^c Full apoptosis variants are identified as *Mna*, e.g., M1a for the full-apoptosis version of “Lopez Embedded”.

Table 2.2: Summary of models in EARM 2.0.

represent genuine mathematical errors or merely transcription errors made in the process of converting computer models to text. Even when we used the ODE networks as published we found cases in which we were unable to reproduce the results described in the figures. Our own previous work was not entirely free of this problem: we could not reproduce the simulation results in Figure 11 of Albeck et al. (Albeck et al., 2008b) without access to MATLAB source code that was inadvertently omitted from the original publication. Our aim is not to criticize these papers but instead to emphasize that the current practice of maintaining different forms of a model for the purpose of simulation, illustration and publication is highly problematic. The lists of equations included as supplementary materials in most modeling papers are particularly troublesome because they exist independently of the simulation model and the two tend to deviate. These problems can be addressed by using electronic formats for model exchange with a single master from which all other versions are derived (Hucka, 2003; Waltemath et al., 2011). As an electronic format for models, PySB complements XML-based formats such as SBML in that macros, modules, and other high-level abstractions make model structure more intelligible than SBML alone. In addition, modeling biochemical processes by reusing previously validated macros eliminates “bookkeeping” errors such as those we identified in published MOMP models. To ensure that the re-instantiated models reproduced the behavior of the originally published versions, we wrote a series of unit tests using the Python modules `unittest` and `nose`. The tests guarantee that the re-instantiated models reproduce validated states despite translation into PySB. Further details on our approach to unit testing can be found in the online documentation (Methods).

The structure and origin of the MOMP models are easier to understand using PySB than the underlying sets of ODEs. This can most easily be seen by comparing the ODE and PySB versions of a model from Chen et al. (C. Chen et al., 2007b) (Figures 2.4B and 2.4C). The original model is relatively simple (only seven ODEs) but understanding the precise mechanism for MOMP requires careful inspection of each equation. By comparison, the PySB model exploits macros to make the mechanisms transparent: single-step catalysis, combinatorial binding, and pore assembly. Many of the 15 MOMP models in EARM 2.0 represent incremental extensions of earlier models (this is particularly true of the models from Howells et al. (2010) and Shen and colleagues (C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008) as

well as the five models from Albeck et al. (2008b; 2008a); Figure 2.4A). The authors of these models proceeded by duplicating ODEs from previous models or papers and then adding new species or reactions as required: for example, the three models of Cui et al. (2008) are derived directly from the “direct” model of Chen et al. (C. Chen et al., 2007b) while the model of Howells et al. (2010) is based on an earlier model from Chen et al. (2007a). However, the process of re-naming of species and variables in the derived models makes it difficult to verify that each variant correctly recapitulates the structure of the original model as claimed. For example, in Cui et al. (2008) the authors state simply that the “direct model” was “mainly based” on the earlier work of Chen et al. (C. Chen et al., 2007b) but we found that there were several important additions and modifications in the derived model, including addition of displacement, synthesis and degradation reactions and a change in the MOMP pore from a Bax tetramer to a Bax dimer. Inspection of the PySB source code for the Cui et al. direct model (Cui et al., 2008) (Figure 2.4D) makes these differences explicit by calling a subroutine for the earlier `chen_febs_direct` model (Figure 2.4C) and adding only the new reactions.

PySB Module Reuse facilitated the process of embedding each of the 15 models of MOMP within the context of receptor-proximate reactions (ligand binding to Bid cleavage) and downstream reactions creating “Full Apoptosis” and “MOMP-only” versions (summarized in Table 2.2; see also Methods). We are currently developing additional apoptosis modules (e.g., alternative topologies for receptor activation and DISC formation) which will soon be part of the EARM repository; other researchers can also “fork” the code on GitHub and contribute their own additions. This should allow a cumulative and distributed approach to model development and comparison.

2.2.5 EMBEDDED TOGETHER: AN UPDATED AND EXPANDED MOMP MODEL

The EARM 2.0 extrinsic apoptosis model incorporating the “Lopez Embedded” MOMP module variant, denoted EARM 2.0-M1a for short (Table 2.2) implements a mathematical interpretation of recent experimental findings from Andrews (Leber et al., 2010; Billen et al., 2008) and Green (Llambi et al., 2011) and differs significantly from previously published models of MOMP (Figure 2.5). Interactions among Bcl-2 family members occur at the mitochondrial membrane rather than in the cytosol (Lovell

et al., 2008), and anti-apoptotic proteins are able to bind both the pore-forming proteins such as Bax and Bak and a larger family of BH3-only Bcl-2 family members, thus serving as dominant-negative effectors (Leber et al., 2010; Billen et al., 2008). This is also consistent with a recent “Unified Model” by Green and coworkers demonstrating both “direct” and “indirect” modes of action by the anti-apoptotic Bcl-2 proteins (Llambi et al., 2011) (Box 2.1). The overlapping binding specificities implied by this model are summarized in a `bind_table` call that includes the key effector for extrinsic apoptosis (Bid), two BH3-only sensitizers (Bad and Noxa), two pore-forming effectors (Bax and Bak) and three anti-apoptotic proteins (Bcl-2, Bcl-XL, and Mcl-1) along with affinity data obtained from *in vitro* experiments (Certo et al., 2006; Willis et al., 2005) (Figure 2.5A). There is some doubt about whether peptide-based affinity measurements are directly relevant to protein-protein interactions occurring on the membranes of living cells, and the `bind_table` macro makes it straightforward to experiment with different values (Figure 2.2C). EARM 2.0-Mia also assumes auto-activation of Bax (and Bak), which has been demonstrated in multiple experimental contexts (Tan et al., 2006; Gavathiotis et al., 2010).

Our previously published EARM 1.0-1.4 models (Albeck et al., 2008b; Albeck et al., 2008a; Gaudet et al., 2012; Spencer et al., 2009; Aldridge et al., 2011) assumed that the all-or-none quality of MOMP arose from the ability of Bcl-2 to bind Bax monomers, dimers and tetramers (Albeck et al., 2008b). However, subsequent immuno-precipitation experiments failed to support the existence of such higher-order hetero-oligomers (Kim et al., 2009). To determine whether the updated reaction topology in EARM 2.0-Mia can reproduce MOMP dynamics measured in single TRAIL-treated HeLa cells using Förster resonance energy transfer (FRET) reporter proteins (Spencer et al., 2009) we fitted it to data using the simulated annealing algorithm in SciPy (Methods). We found that EARM 2.0-Mia had as good a fit to data as previous models (Figure 2.6) and we therefore judge it to be superior to our earlier EARM 1.0 model based simply on better correspondence with prior knowledge. The fitting exercise also demonstrated that Python numerical tools can efficiently simulate and calibrate PySB models (parameter estimation functions are included in the EARM 2.0 Python package).




Figure 2.5: Representations of the Bcl-2 interaction topology in the EARM 2.0-M1b (“Lopez Embedded”) model. The Bcl-2 interaction model consists of five basic mechanistic elements or “motifs”—tBid/Bax/BclxL translocation, activation of Bax and Bak by tBid, Bcl-2 family binding, Bax/Bak autoactivation, and pore assembly—that are highlighted in (A), (B) and (C) according to the colors in the legend.

(A) PySB source code for the model, edited for brevity.

(B) Simplified, manually drawn representation.

(C) The full reaction network, generated from the PySB model using the PySB `render_reactions` tool. Rectangles represent species, circles represent reactions, lines represent reactions with the solid arrowhead representing the nominal forward direction and the empty arrowhead (for reversible reactions only) representing the reverse direction. Catalytic reactions are depicted with a boxed arrow pointing from the catalyst to the reaction circle (species for enzyme-substrate complexes are omitted for clarity).

(D) Kappa contact map, which shows the superset of all possible bonds between monomers calculated by static analysis (Danos et al., 2008). The contact map was computed using Kappa’s `compLx` tool accessed through the PySB Kappa wrapper library. Rectangles represent monomers, circles represent sites and lines represent bonds.




Figure 2.6: Simulated annealing fits for EARM 2.0-M1a for three experimentally measured protein signals IC-RP/tBid (**A**), IMS-RP/Smac release (**B**), and EC-RP/PARP cleavage (**C**). Gray lines indicate the experimental data with error bars indicating the s.d. In the case of (**B**), the gray line denotes the mean time of death T_D^{Exp} , used to align the trajectories (**Methods**). The orange curves show the simulated trajectories using nominal parameter values; the green curves show the simulated trajectories after model fitting. The objective function for fitting was as described in **Methods**.

2.3 DISCUSSION

In this paper we describe the development and use of PySB, a framework for creating, managing and analyzing biological models in which models are full-fledged Python programs. PySB modules and macros generate BioNetGen Language or Kappa rules, which can be converted into mathematical equations. This hierarchical process is analogous to the way in which programs are written in a high-level language such as C++ and converted into microprocessor code by the compiler. This complexity is hidden from PySB users who work with macro libraries encoding biochemical actions such as “catalyze,” “bind,” “assemble,” etc. The advantages of a high-level abstraction layer include greater transparency and intelligibility, a reduction in implementation errors and a dramatic increase in the ability to compare and reuse

previous work. Each level of representation remains accessible for analysis and there are no explicitly black-box steps: the Python model code reveals model ontogeny, structure and approach to mechanism; the BNGL or Kappa rules generated by PySB support agent-based simulation and static analysis; and equations enable numerical integration, sensitivity analysis, steady-state analysis, etc. Use of familiar but powerful programming concepts in PySB models such as abstraction, composition, modularity, inheritance, and polymorphism make it possible to create variant models from pre-existing models across several axes of variation and build new models from previously tested elements. We expect these features of PySB to facilitate collaborative model development and evaluation.

PySB draws on well-established practices in the open-source programming community for model documentation and sharing. Because PySB models are programs they can be tracked and shared using the powerful tools developed for distributed, open-source software development (e.g., all the models in this paper are available, with documentation, at GitHub; see [Methods](#)). It is simple to update models online, highlight differences with previous work and divide development among multiple individuals and research teams. Finally, PySB can be used as a general-purpose modeling tool because it interoperates with diverse scientific applications written in Python (e.g. NumPy, SciPy, SymPy, and Matplotlib). Unlike conventional all-in-one programs PySB itself tackles only certain steps in the modeling process, relying on interoperability with programs developed and maintained by others to create a full-fledged solution. A benefit of this approach is that improvements in any of these programs accrue directly to users of PySB.

The power of PySB derives, in part, from its ability to encode recurrent biochemical patterns in reusable macros ([Figure 2.2](#)) and to divide complex networks into modules that are defined once and called when needed ([Figures 2.3B](#) and [Figure 2.4C](#)). By eliminating re-implementation, macros and modules separate fundamental mechanistic concepts from implementation details and thereby make clear the purpose and origins of specific model features ([Mallavarapu et al., 2009](#); [Gnad et al., 2012](#); [Pedersen and Plotkin, 2008](#); [Mirschel et al., 2009](#)). The ability of real biological networks to be meaningfully decomposed into functional modules is highly context dependent and a matter of controversy ([Del Vecchio et al., 2008](#)) but there is no requirement that modules in PySB correspond to modules in a biological or “black-box” engineering sense: the full reaction network is always accessible without simplification. Instead, PySB

modules are defined according to flexible and convenient organizational boundaries, keeping open the possibility for crosstalk and emergent interactions with other modules. This style of modularity follows the open-ended approach of *little b* (Mallavarapu et al., 2009) and differs from ProMoT, in which modules interact only through previously designated molecular species (Mirschel et al., 2009). In general choosing the right boundaries for a module, whether a software program or a biological model, is a matter of art and practical experience. In the models of extrinsic apoptosis analyzed in this paper, reactions governing MOMP are a good candidate for modularization because they largely take place in a discrete compartment (the mitochondrial membrane) and focus on reactions among Bcl-2 proteins.

We have found that PySB naturally supports a hierarchy of modeling concepts (Figure 2.7A). At the top of this hierarchy are the models themselves, which represent a specific hypothesis about the topology and activity of a biological system or network; at the bottom of the hierarchy are specific mathematical equations (e.g. ODEs). Typical approaches to modeling proceed by directly rendering the hypothesis in equations, making it difficult to discern the assumptions implicit in the process of mathematical translation (Figure 2.7B). Rule-based approaches represent an intermediate level of abstraction in that they enumerate local interactions between proteins in a way that is less explicit than equations (Figure 2.7C). PySB adds an additional layer of abstraction in that the user works with macros and functions (Figure 2.7A; see also Figure 2.2). Sets of macros are then grouped into reusable subroutines that implement small mechanistic “motifs” corresponding roughly to a sentence in a word model, such as “tBid activates Bax and Bak.” Such “motifs” are then composed into modules, and modules into models. Constructed in this fashion, a set of variant models forms a “web” of intertwined elements that is largely self-documenting.

2.3.1 PYSB AS A SECOND-GENERATION APPROACH

PySB is not the first attempt to create a high-level language for modeling biochemistry and was inspired by ProMoT and *little b*, both of which represent models as LISP programs (Mirschel et al., 2009; Mallavarapu et al., 2009). However, as described above, these tools had limited or no support for rule-based modeling. PySB is based on the much more familiar Python language and is interoperable with BNGL and Kappa.

The rule-based modeling community is also developing tools for managing complex models. For exam-




Figure 2.7: A conceptual hierarchy for model building with PySB.

(A) *Models* can be decomposed into a small number of separable *modules*; modules can in turn be decomposed into smaller recurrent mechanistic elements or *motifs*. Motifs are collections of *macro* calls (e.g., the motif “tBid activates effectors” consists of two *cat-*
a_lyze calls; Figure 2.5A). Each macro expands into multiple rules (Figure 2.2) and rules generate equations (Figure 2.1B).
(B) ODE-based approaches to model construction proceed directly from a complex biological system to its mathematical representa-
tion, while (C) rule-based approaches provide one abstraction layer above ODEs.

ple, MetaKappa targets redundancy in models having related molecular species and partially overlapping functional characteristics (e.g. a set of mutants or isoforms of a single protein) (Danos et al., 2009). The Language for Biological Systems (LBS) is another approach in which rules are combined with methods for constructing parameterized modules (Pedersen and Plotkin, 2008). MetaKappa and LBS are examples of domain-specific languages (DSLs) for high-level biological modeling, whereas PySB supports high-level modeling through the structured programming features of Python (Box 2.2). Through Python, PySB provides substantial flexibility in organizing models. A potential drawback of this flexibility is that static analysis of models may be more challenging to implement than for DSLs such as MetaKappa or LBS. For the time being, available static analysis algorithms can be applied to the rules generated from PySB models (Danos et al., 2008; Feret et al., 2009). We expect the strengths and weaknesses of these different approaches to become clearer as high-level languages are more widely adopted by the modeling community.

Graphical tools represent an alternative approach to making biological models easier to understand. CellDesigner (Kitano et al., 2005) and BioUML (Kutumova et al., 2012) employ visual interfaces and graphical languages for model creation, interpretation, and revision. Our experience has been that purely visual approaches to model creation do not scale well with model size and are ambiguous with respect to the underlying biochemical mechanisms. This parallels experience in the software engineering commu-

nity with UML (Unified Modeling Language) diagrams: diagrams are a helpful adjunct to programmatic abstractions but they are not a replacement. This implies that visualizations should be created from an underlying computable representation and not the other way around; this is the approach taken by the *rxncon* software, which uses specialized tables to describe reactions in a way that specifically supports multiple visualizations (Tiger et al., 2012). However, to be most useful, the computable representation must be made as transparent as possible: approaches that rely on generated visualizations for transparency are hard to modify, since changes must be made to the underlying (non-transparent) representation. While visualization remains important for communicating and describing models (Chylek et al., 2011), PySB follows the common software engineering paradigm in which programmatic abstraction serves as the principal tool for managing complexity and visualization serves to illustrate specific properties of a system (Danos et al., 2008). By way of illustration, we show how a reaction network (Figure 2.5C) and a Kappa contact map (Figure 2.5D) can be generated from EARM 2.0-M1b (the MOMP-module-only variant of M1a) and a species graph can be generated for the full EARM 2.0-M1a model (Supplementary Figure S2). These visualizations proved useful in debugging models during the preparation of this manuscript but it was the PySB model code that supported straightforward revisions (Figure 2.5A). The code not only contains more information than the visualizations, but with only 10 macro calls falling into 5 “motifs” (color-coded regions) it compares favorably with even a simplified, hand-drawn “cartoon” representation (Figure 2.5B) in terms of intelligibility.

The advantages of using programmatic methods for modeling biological pathways are not necessarily evident from simple examples in which the underlying equations are self-explanatory (e.g. the “Hello World” model in Figure 2.1A). This has led some to dispute the value of such methods and to argue that direct modeling in equations is superior. However, as models become more realistic and complex, equations rapidly become difficult to understand and errors creep in, a problem that was evident with the published MOMP models we re-instantiated (Figure 2.4). The advantages of programmatic abstraction also become more evident when a model must be revised or shared, particularly if the original developers have moved on to something new. This paper shows how PySB, in combination with BNGL and Kappa, goes a long way towards addressing this problem.

2.3.2 THE EARM 2.0 MODELS OF EXTRINSIC CELL DEATH

One aim of this paper was to create a new model of extrinsic apoptosis that incorporated the latest thinking on the biochemistry of MOMP while facilitating comparison with previously published models. EARM 2.0 includes 15 different models for the reactions controlling MOMP, 12 of which were previously published in five papers and three of which were novel. In general, previously published models of MOMP do not explore Bcl-2 biochemistry in the context of a complete receptor-to-effector caspase network. Such models also simplify the biology of Bcl-2 proteins, representing only a subset of the family members. We overcame these limitations by modularizing the extrinsic apoptosis pathway and using composition to embed different MOMP modules within a larger network. Use of the `bind_table` macro (for modeling interactions among members of a multi-protein family) made it possible to efficiently encode the differential binding affinities of many Bcl-2 proteins for each other. While the molecular interactions included in our EARM 2.0 models are not comprehensive, the extensible nature of the PySB representation makes it simple to add additional mechanistic details and species in future work.

One important feature of apoptosis is that MOMP pores do not form until a pro-death stimulus has been applied at a sufficient concentration for several hours; however, once pores form cells die quickly (Albeck et al., 2008b; Goldstein et al., 2000; Rehm et al., 2002). Our original explanation for the all-or-none regulation of MOMP was that an anti-apoptotic protein such as Bcl-2 or Bcl-XL binds to successively larger Bax and Bak oligomers, thereby creating a cooperative inverse relationship between Bcl-2 and pore levels. This idea has not been borne out by experiments and instead it appears that the kinetic properties of MOMP must arise from the dual affinity of anti-apoptotic proteins for both BH₃-only proteins and pore-forming effectors (Llambi et al., 2011; Billen et al., 2008). We find that an “embedded together” model (EARM 2.0-M1a) incorporating this revised thinking about Bcl-2 family proteins (Leber et al., 2010) can reproduce the dynamics of MOMP as measured in single cells using live-cell imaging (Spencer et al., 2009). An additional round of experimental and computational “model discrimination” studies is needed to show whether EARM 2.0-M1a is indeed superior to previous models but we prefer it simply on the basis of its faithful recapitulation of current knowledge.

2.3.3 PySB AS A MEANS TO INCREMENTAL AND COLLABORATIVE MODEL DEVELOPMENT

One of the key aims for PySB and also for BNGL, Kappa and related meta-languages, is to promote distributed, incremental and collaborative approaches to modeling. There are both technical and conceptual challenges that must be addressed for this to be successful. A significant technical hurdle in model integration and reuse is the need for a standard nomenclature for model species. The SBML community's MIRIAM standard is an essential resource in this regard ([Le Novère et al., 2005](#)), but the way in which rule-based models represent species and complexes will demand a modified approach to annotation. As a first step we have implemented a basic annotation capability in PySB (see [Supplementary Figure S3](#)) based on MIRIAM-style subject-object-predicate triples that should help resolve naming ambiguities.

However, the fundamental challenge for integrating and reusing models of disparate biological processes remains the fact that biological models remain “fit-to-purpose,” focused on addressing specific biological problems or contexts ([Krakauer et al., 2011](#)). PySB does not prescribe “universal” approaches to representing biological components or processes but instead makes fit-to-purpose modeling more transparent and manageable through the use of both hierarchical abstractions ([Figure 2.7](#)) and tools for documenting, testing, and tracking models drawn from software engineering. In the short-term, these features should allow communities of biologists working on related biological problems to work in parallel toward shared goals; in the longer-term, real-world experimentation with approaches for collaborative modeling should yield best practices for building broadly reusable models.

The fact that PySB models are programs allows us to exploit the tools and social conventions of the open-source software development community for distributed model development. In open-source software, derivative or variant branches of a source tree can be spun off and then merged into the principal source tree if desired. Version control systems such as Git allow this process to be managed and visualized. PySB models shared via GitHub can have both private and public branches that preserve the integrity of ongoing model development while allowing for external contributions. Multiple groups can develop derivative models with confidence that the relationships among variants can be tracked and managed. SBML versions of PySB models can also be deposited in [biomodels.net](#), supporting current procedures for indexing, citation, and search. Software tests can be written to ensure that models and modules be-

have as documented. As models get larger and the scope of the underlying biology exceeds the expertise of a single modeling team, tools such as PySB will be needed to create reusable, shareable and transparent biological models in a distributed manner—a major goal of a systems-level program of biological discovery.

2.4 METHODS

2.4.1 PySB CODE AND DOCUMENTATION

PySB is freely available under an open-source BSD license. Links to the GitHub source code repository as well as documentation and other didactic materials are available at <http://pysb.org>. The EARM 2.0 models and associated documentation are available along with the data used for model calibration at <http://sorgerlab.github.com/earm>. An SBML version of the EARM 2.0-M1a model is included in the [Supplementary information](#) of the paper.

2.4.2 PySB SYNTAX

PySB implements a syntax for rules based on that of BNGL using classes and overloaded operators from the Python language, leveraging Python as the language parser. Rule definitions are built up using patterns (represented internally as instances of the classes `MonomerPattern`, `ComplexPattern`, and `ReactionPattern`) that specify both a rule's preconditions and its effects. Site and state conditions for monomers are specified using an overloaded `__call__` method for the `Monomer` class, which takes site and state conditions as keyword arguments. For example, if `L` is an instance of `Monomer`, `L(s='P')` specifies `L` with site `s` having state `P`. This use of the `__call__` method, along with the overloaded operators `+`, `%`, `>>` and `<>`, allow rules to be specified using a syntax that parallels that of BNGL and Kappa ([Figure 2.1A](#)). See also the “PySB Syntax” section of the [Supplementary information](#).

2.4.3 COMPONENT SELF-EXPORT

By default, when a model component is constructed, it is added to the current instance of the `Model` class; a variable referring to the newly created component (with a name matching the name passed to the

component constructor) is also inserted into the global namespace. This eliminates the need to retain references to newly created objects and explicitly add them to the model. The “self-export” functionality is managed by the PySB core class `SelfExporter`, which identifies the module in which the current `Model` instance was declared and adds global variables for components to that namespace. If alternative approaches to component and model management are desired, the self-export feature can be disabled by adding the keyword argument `_export=False` to each component constructor. See also the “PySB Syntax” section of the [Supplementary information](#).

2.4.4 ODE INTEGRATION

PySB generates the reaction network through an external call to BioNetGen and extracts the network graph by parsing the resulting `.net` file. The network graph is then used to build up the right-hand sides of the ODEs as Sympy symbolic math expressions containing the appropriate rate constants, mass action terms and stoichiometric coefficients. If a C compiler is available, the right-hand side function is implemented in C using `scipy.weave.inline`; otherwise the right-hand side function is evaluated as a Python expression. A reference to the right-hand side function is passed to `scipy.integrate.ode`, a generic wrapper library for ODE integration routines; for this work we used the FORTRAN integrator VODE ([Brown et al., 1989](#)).

2.4.5 CONVERSION OF PUBLISHED MODELS TO PySB

For the models lacking an electronic version ([C. Chen et al., 2007a](#); [C. Chen et al., 2007b](#); [Cui et al., 2008](#); [Howells et al., 2010](#)), the ODEs generated by the PySB versions were manually validated against the ODEs listed in the original publications. In the case of the three models with errors in the published ODEs ([Cui et al., 2008](#)) (see [Supplementary note](#)) the PySB version was written to generate the ODEs corresponding to the described reaction scheme without these errors. For the models for which we had access to the original MATLAB code ([Albeck et al., 2008b](#)), the PySB versions were also programmatically validated against the output from the published versions.

2.4.6 MODULARIZATION OF MOMP MODELS

Each MOMP-only model (Table 2.2) was written to have the addition of tBid as its most upstream event, and the release of cytochrome *c* and Smac as its most downstream event. In some cases these default boundaries did not match the boundaries for the MOMP module in the original publications: Albeck et al. had the addition of active caspase-8 as the most upstream event (Albeck et al., 2008b), while the Shen/Howells group models had Bax oligomerization (rather than Smac release) as the most downstream event (C. Chen et al., 2007a; C. Chen et al., 2007b; Cui et al., 2008; Howells et al., 2010). In these cases, the networks of the original models were modified to achieve consistent boundaries across modules. The boundaries of the original models can nevertheless be reproduced through the use of optional arguments to the module subroutines that add or remove reactions as appropriate.

2.4.7 SIMULATION AND PARAMETER ESTIMATION

Simulations were carried out using the VODE integrator via the SciPy library using Newton's method for root evaluations and the backward differentiation formula (BDF) integration method. Absolute and relative tolerances were set to 10^{-5} . Parameter estimation was performed using the simulated annealing routine implemented in `scipy.optimize.anneal` with an appropriately defined objective function (described below). Nominal values for rate constants for the DISC and PARP modules were set to their published values in EARM 1.0; rate constants for the MOMP module were drawn from Certo et al. (2006), Willis et al. (2005), or set to values from similar rates from EARM 1.0. During the annealing process all rate constraints were allowed to vary two orders of magnitude above and below their nominal values (i.e., 0.01x – 100x); initial protein concentrations were held fixed and not estimated.

Trajectories for the initiator caspase reporter protein (IC-RP), mitochondrial inter-membrane space reporter protein (IMS-RP) and effector caspase reporter protein (EC-RP) were used from previously published data (Spencer et al., 2009). In the model, truncated Bid, cytosolic Smac, and cleaved PARP were fit to the data for IC-RP, IMS-RP, and EC-RP, respectively. IMS-RP data from ten cells indicated an average MOMP time of $9810 \text{ s} \pm 2690 \text{ s}$ after the exposure of the cells to ligand. The IC-RP and EC-RP signals were normalized and aligned to this MOMP time to yield an average trajectory for each. The objective

function used to calculate model fit was the sum of component functions for each of the experimental observables as follows:

$$Obj_{Total}(\vec{k}) = Obj_{ICRP}(\vec{k}) + Obj_{IMSRP}(\vec{k}) + Obj_{ECRP}(\vec{k})$$

where $\vec{k} = k_1, k_2, \dots, k_n$ are the rate parameters for the system of ODEs, and $Obj_{ICRP}(\vec{k})$, $Obj_{IMSRP}(\vec{k})$, and $Obj_{ECRP}(\vec{k})$ are the objective functions for corresponding observables. These objective functions were defined as follows:

$$Obj_{ICRP}(\vec{k}) = \frac{1}{N} \sum_{t=1}^N \frac{\left[\frac{tBid_t(\vec{k})}{Bid_o} - \langle ICRP_t \rangle \right]^2}{2 \cdot Var(ICRP_t)}$$

where $N = 112$ is the number of experimental timepoints, $tBid_t(\vec{k})$ is the simulated value of truncated Bid at time t (with the given parameters, k), Bid_o is the total amount of Bid in the simulation, and $ICRP_t$ is the distribution of MOMP-aligned experimental IC-RP trajectories at time t . The objective function for EC-RP was defined similarly, but with cleaved PARP in place of truncated Bid and EC-RP in place of IC-RP:

$$Obj_{ECRP}(\vec{k}) = \frac{1}{N} \sum_{t=1}^N \frac{\left[\frac{cPARP_t(\vec{k})}{PARP_o} - \langle ECRP_t \rangle \right]^2}{2 \cdot Var(ECRP_t)}$$

Data from previous experiments showed that release of IMS-RP from mitochondria due to MOMP is both rapid and complete under all tested experimental conditions, typically reaching completion within 1-3 minutes (Albeck et al., 2008a). To ensure that the model could reproduce these kinetics, we used an alternative objective function for the cSmac/IMS-RP signal to consider the onset (T_D) and duration (T_S) of MOMP:

$$Obj_{IMSRP}(\vec{k}) = \frac{1}{3} \left[\frac{\left(T_D^{Sim} - \langle T_D^{Exp} \rangle \right)^2}{2 \cdot Var(T_D^{Exp})} + \frac{\left(T_S^{Sim} - \langle T_S^{Exp} \rangle \right)^2}{2 \cdot Var(T_S^{Exp})} + \frac{\left(\frac{cSmac_N}{Smac_o} - 1 \right)^2}{2 \cdot 10^{-6}} \right]$$

T_D^{Exp} is the experimental “death time,” recorded as the time of MOMP as measured by IMS-RP release; T_S^{Exp} is the experimental “switching time,” defined as the time required for IMS-RP to be released, estimated to have a mean value of 180 sec and a variance of 3600 sec². The corresponding simulated values were defined as

$$T_D^{sim} = \frac{T_{90} + T_{10}}{2}$$

$$T_S^{sim} = T_{90} - T_{10}$$

where T_{90} and T_{10} are the times taken by the simulation to reach 90% and 10% of maximum Smac release, respectively. The final term in $Obj_{IMSRP}(\vec{k})$ constrains the simulated Smac trajectory to achieve approximately 100% release: $cSmac_N$ is the final simulated value of cytosolic Smac, and $Smac_o$ is total Smac.

ACKNOWLEDGEMENTS

The authors thank A. Keating and E. Fire for assistance in the early stages of this project and W. Fontana, J. Gunawardena and A. Mallavarapu for help and advice. This work was supported by NIH grants P50-GM68762 and P01-CA139980. CFL was supported by The Harold Whitworth Pierce Charitable Trust (through The Medical Foundation) and NIH Transition to Independence grant K22-CA151918.

References

- Albeck, J. G., J. M. Burke, B. B. Aldridge, M. Zhang, D. A. Lauffenburger, and P. K. Sorger (2008a). “Quantitative analysis of pathways controlling extrinsic apoptosis in single cells”. *Molecular Cell* 30.1 (Apr. 2008), pp. 11–25.
- Albeck, J. G., J. M. Burke, S. L. Spencer, D. A. Lauffenburger, and P. K. Sorger (2008b). “Modeling a snap-action, variable-delay switch controlling extrinsic cell death”. *PLoS Biology* 6.12 (Dec. 2008), pp. 2831–2852.
- Aldridge, B. B., S. Gaudet, D. A. Lauffenburger, and P. K. Sorger (2011). “Lyapunov exponents and phase diagrams reveal multi-factorial control over TRAIL-induced apoptosis”. *Molecular Systems Biology* 7, p. 553.
- Bachman, J. A. and P. Sorger (2011). “New approaches to modeling complex biochemistry.” *Nature Methods* 8.2 (Feb. 2011), pp. 130–131.
- Bagci, E., Y. Vodovotz, T. Billiar, G. Ermentrout, and I. Bahar (2006). “Bistability in apoptosis: roles of bax, bcl-2, and mitochondrial permeability transition pores”. *Biophysical Journal* 90.5, pp. 1546–1559.
- Batchelor, E., C. S. Mock, I. Bhan, A. Loewer, and G. Lahav (2008). “Recurrent initiation: a mechanism for triggering p53 pulses in response to DNA damage.” *Molecular Cell* 30.3 (May 2008), pp. 277–289.
- Bentele, M., I. Lavrik, M. Ulrich, S. Stösser, D. W. Heermann, H. Kalthoff, P. H. Krammer, and R. Eils (2004). “Mathematical modeling reveals threshold mechanism in CD95-induced apoptosis”. *The Journal of Cell Biology* 166.6 (Sept. 2004), pp. 839–851.
- Billen, L. P., C. L. Kokoski, J. F. Lovell, B. Leber, and D. W. Andrews (2008). “Bcl-XL inhibits membrane permeabilization by competing with Bax”. *PLoS Biology* 6.6 (June 2008), e147.
- Blinov, M. L., J. R. Faeder, B. Goldstein, and W. S. Hlavacek (2006). “A network model of early events in epidermal growth factor receptor signaling that accounts for combinatorial complexity.” *Bio Systems* 83.2-3 (Jan. 2006), pp. 136–151.
- Brown, P., G. Byrne, and A. C. Hindmarsh (1989). “VODE, A variable coefficient ODE solver”. *SIAM J Sci Stat Comput* 10, pp. 1038–1051.
- Certo, M., V. Del Gaizo Moore, M. Nishino, G. Wei, S. Korsmeyer, S. A. Armstrong, and A. Letai (2006). “Mitochondria primed by death signals determine cellular addiction to antiapoptotic BCL-2 family members”. *Cancer Cell* 9.5 (May 2006), pp. 351–365.
- Chen, C., J. Cui, H. Lu, R. Wang, S. Zhang, and P. Shen (2007a). “Modeling of the role of a Bax-activation switch in the mitochondrial apoptosis decision”. *Biophysical Journal* 92.12 (June 2007), pp. 4304–4315.

- Chen, C., J. Cui, W. Zhang, and P. Shen (2007b). "Robustness analysis identifies the plausible model of the Bcl-2 apoptotic switch". *FEBS letters* 581.26 (Oct. 2007), pp. 5143–5150.
- Chen, L. et al. (2005). "Differential Targeting of Prosurvival Bcl-2 Proteins by Their BH₃-Only Ligands Allows Complementary Apoptotic Function". *Molecular Cell* 17.3 (Apr. 2005), pp. 393–403.
- Chen, W. W., B. Schoeberl, P. J. Jasper, M. Niepel, U. B. Nielsen, D. A. Lauffenburger, and P. K. Sorger (2009). "Input-output behavior of ErbB signaling pathways as revealed by a mass action model trained against dynamic data". *Molecular Systems Biology* 5, p. 239.
- Chylek, L. A. et al. (2011). "Guidelines for visualizing and annotating rule-based models." *Molecular BioSystems* 7.10 (Oct. 2011), pp. 2779–2795.
- Cui, J., C. Chen, H. Lu, T. Sun, and P. Shen (2008). "Two independent positive feedbacks and bistability in the Bcl-2 apoptotic switch". *PLoS ONE* 3.1, e1469.
- Danos, V., J. Feret, W. Fontana, and J. Krivine (2007a). "Scalable Simulation of Cellular Signaling Networks". *LECTURE NOTES IN COMPUTER SCIENCE*.
- Danos, V. (2009). "Agile Modelling of Cellular Signalling (Invited Paper)". *Electronic Notes in Theoretical Computer Science* 229.4, pp. 3–10.
- Danos, V., J. Feret, W. Fontana, R. Harmer, and J. Krivine (2009). "Rule-Based Modelling and Model Perturbation". In: *Transactions on Computational Systems Biology XI*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 116–137.
- Danos, V., J. Feret, W. Fontana, R. Harmer, and J. Krivine (2007b). "Rule-Based Modelling of Cellular Signalling". In: *CONCUR 2007 – Concurrency Theory*. Berlin, Heidelberg: Springer Berlin Heidelberg, Sept. 2007, pp. 17–41.
- Danos, V., J. Feret, W. Fontana, and J. Krivine (2008). *Abstract Interpretation of Cellular Signalling Networks*. Springer Berlin Heidelberg. Berlin, Heidelberg.
- Debartolo, J., S. Dutta, L. Reich, and A. E. Keating (2012). "Predictive bcl-2 family binding models rooted in experiment or structure." *Journal of Molecular Biology* 422.1 (Sept. 2012), pp. 124–144.
- Deeds, E. J., J. Krivine, J. Feret, V. Danos, and W. Fontana (2012). "Combinatorial complexity and compositional drift in protein interaction networks." *PLoS ONE* 7.3, e32032.
- Del Vecchio, D., A. J. Ninfa, and E. D. Sontag (2008). "Modular cell biology: retroactivity and insulation." *Molecular Systems Biology* 4, p. 161.
- Dussman, H., M. Rehm, C. G. Concannon, S. Anguissola, M. Wurstle, S. Kacmar, P. Voiler, H. J. Huber, and J. H. M. Prehn (2009). "Single-cell quantification of Bax activation and mathematical modelling suggest pore formation on minimal mitochondrial Bax accumulation". *Cell Death and Differentiation* (Nov. 2009), pp. 1–13.
- Faeder, J. R., M. L. Blinov, and W. S. Hlavacek (2009). "Rule-based modeling of biochemical systems with BioNetGen." *Methods in molecular biology (Clifton, NJ)* 500. Chapter 5, pp. 113–167.

- Feret, J., V. Danos, J. Krivine, R. Harmer, and W. Fontana (2009). "Internal coarse-graining of molecular systems". *Proceedings of the National Academy of Sciences of the United States of America* 106.16 (Apr. 2009), pp. 6453–6458.
- Fire, E., S. V. Gullá, R. A. Grant, and A. E. Keating (2010). "Mcl-1-Bim complexes accommodate surprising point mutations via minor structural changes." *Protein Science* 19.3 (Mar. 2010), pp. 507–519.
- Fricker, N., J. Beaudouin, P. Richter, R. Eils, P. H. Krammer, and I. N. Lavrik (2010). "Model-based dissection of CD95 signaling dynamics reveals both a pro- and antiapoptotic role of c-FLIPL." *The Journal of Cell Biology* 190.3 (Aug. 2010), pp. 377–389.
- Gansner, E. R. and S. C. North (2000). "An open graph visualization system and its applications to software engineering". *Software Practice and Experience*.
- Gaudet, S., S. L. Spencer, W. W. Chen, and P. K. Sorger (2012). "Exploring the contextual sensitivity of factors that determine cell-to-cell variability in receptor-mediated apoptosis." *PLoS Computational Biology* 8.4, e1002482.
- Gavathiotis, E., D. E. Reyna, M. L. Davis, G. H. Bird, and L. D. Walensky (2010). "BH₃-Triggered Structural Reorganization Drives the Activation of Proapoptotic BAX". *Molecular Cell* 40.3 (Nov. 2010), pp. 481–492.
- Gnad, F., J. Estrada, and J. Gunawardena (2012). "Proteus: a web-based, context-specific modelling tool for molecular networks." *Bioinformatics (Oxford, England)* 28.9 (May 2012), pp. 1284–1286.
- Goldstein, J. C., R. M. Kluck, and D. R. Green (2000). "A single cell analysis of apoptosis. Ordering the apoptotic phenotype". *Annals of the New York Academy of Sciences* 926, pp. 132–141.
- Hlavacek, W. S. (2009). "How to deal with large models?" *Molecular Systems Biology* 5.1, p. 240.
- Hlavacek, W. S., J. R. Faeder, M. L. Blinov, R. G. Posner, M. Hucka, and W. Fontana (2006). "Rules for modeling signal-transduction systems." *Science's STKE : signal transduction knowledge environment* 2006.344 (July 2006), re6–re6.
- Hoops, S. et al. (2006). "COPASI—a COMplex PATHway SIMulator." *Bioinformatics (Oxford, England)* 22.24 (Dec. 2006), pp. 3067–3074.
- Howells, C. C., W. T. Baumann, D. C. Samuels, and C. V. Finkielstein (2010). "The Bcl-2-associated death promoter (BAD) lowers the threshold at which the Bcl-2-interacting domain death agonist (BID) triggers mitochondria disintegration". *Journal of Theoretical Biology* (Dec. 2010).
- Huber, H. J., H. Dussmann, S. M. Kilbride, M. Rehm, and J. H. M. Prehn (2011). "Glucose metabolism determines resistance of cancer cells to bioenergetic crisis after cytochrome-c release." *Molecular Systems Biology* 7 (Mar. 2011), p. 470.
- Hucka, M. (2003). "The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models". *Bioinformatics (Oxford, England)* 19.4 (Mar. 2003), pp. 524–531.

- Hunter, J. D. (2007). "Matplotlib: A 2D Graphics Environment". *Computing in Science & Engineering* 9.3 (May 2007), pp. 90–95.
- Kim, H., M. Rafiuddin-Shah, H.-C. Tu, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2006). "Hierarchical regulation of mitochondrion-dependent apoptosis by BCL-2 subfamilies". *Nature Cell Biology* 8.12 (Dec. 2006), pp. 1348–1358.
- Kim, H., H.-C. Tu, D. Ren, O. Takeuchi, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2009). "Stepwise activation of BAX and BAK by tBID, BIM, and PUMA initiates mitochondrial apoptosis". *Molecular Cell* 36.3 (Nov. 2009), pp. 487–499.
- Kitano, H., A. Funahashi, Y. Matsuoka, and K. Oda (2005). "Using process diagrams for the graphical representation of biological networks." *Nature Biotechnology* 23.8 (Aug. 2005), pp. 961–966.
- Kleiman, L. B., T. Maiwald, H. Conzelmann, D. A. Lauffenburger, and P. K. Sorger (2011). "Rapid phospho-turnover by receptor tyrosine kinases impacts downstream signaling and drug binding." *Molecular Cell* 43.5 (Sept. 2011), pp. 723–737.
- Knuth, D. E. (1984). "Literate Programming". *The Computer Journal* 27.2 (Jan. 1984), pp. 97–111.
- Krakauer, D. C., J. P. Collins, D. Erwin, J. C. Flack, W. Fontana, M. D. Laubichler, S. J. Prohaska, G. B. West, and P. F. Stadler (2011). "The challenges and scope of theoretical biology." *Journal of Theoretical Biology* 276.1 (May 2011), pp. 269–276.
- Kutumova, E. O., I. N. Kiselev, R. N. Sharipov, I. N. Lavrik, and F. A. Kolpakov (2012). "A modular model of the apoptosis machinery." *Advances in experimental medicine and biology* 736, pp. 235–245.
- Kuwana, T., L. Bouchier-Hayes, J. E. Chipuk, C. Bonzon, B. A. Sullivan, D. R. Green, and D. D. Newmeyer (2005). "BH3 domains of BH3-only proteins differentially regulate Bax-mediated mitochondrial membrane permeabilization both directly and indirectly." *Molecular Cell* 17.4 (Feb. 2005), pp. 525–535.
- Le Novère, N. et al. (2005). "Minimum information requested in the annotation of biochemical models (MIRIAM)." *Nature Biotechnology* 23.12 (Dec. 2005), pp. 1509–1515.
- Leber, B., J. Lin, and D. W. Andrews (2010). "Still embedded together binding to membranes regulates Bcl-2 protein interactions." *Oncogene* 29.38 (Sept. 2010), pp. 5221–5230.
- Legewie, S., N. Blüthgen, and H. Herzel (2006). "Mathematical Modeling Identifies Inhibitors of Apoptosis as Mediators of Positive Feedback and Bistability". *PLoS Computational Biology* 2.9 (Sept. 2006), e120.
- Letai, A., M. C. Bassik, L. D. Walensky, M. D. Sorcinelli, S. Weiler, and S. J. Korsmeyer (2002). "Distinct BH3 domains either sensitize or activate mitochondrial apoptosis, serving as prototype cancer therapeutics". *Cancer Cell* 2.3 (Sept. 2002), pp. 183–192.
- Llambi, F., T. Moldoveanu, S. W. G. Tait, L. Bouchier-Hayes, J. Temirov, L. L. McCormick, C. P. Dillon, and D. R. Green (2011). "A unified model of mammalian BCL-2 protein family interactions at the mitochondria." *Molecular Cell* 44.4 (Nov. 2011), pp. 517–531.

- Lovell, J. F., L. P. Billen, S. Bindner, A. Shamas-Din, C. Fradin, B. Leber, and D. W. Andrews (2008). "Membrane binding by tBid initiates an ordered series of events culminating in membrane permeabilization by Bax". *Cell* 135.6 (Dec. 2008), pp. 1074–1084.
- Maiwald, T. and J. Timmer (2008). "Dynamical modeling and multi-experiment fitting with Potter'sWheel." *Bioinformatics (Oxford, England)* 24.18 (Sept. 2008), pp. 2037–2043.
- Mallavarapu, A., M. Thomson, B. Ullian, and J. Gunawardena (2009). "Programming with models: modularity and abstraction provide powerful capabilities for systems biology". *Journal of The Royal Society Interface* 6.32 (Mar. 2009), pp. 257–270.
- Martinez-Caballero, S., L. M. Dejean, M. S. Kinnally, K. J. Oh, C. A. Mannella, and K. W. Kinnally (2009). "Assembly of the Mitochondrial Apoptosis-induced Channel, MAC". *The Journal of biological chemistry* 284.18 (May 2009), pp. 12235–12245.
- Mathworks, T. *MATLAB*. Natick, MA: The Mathworks.
- Mérino, D. et al. (2009). "The role of BH3-only protein Bim extends beyond inhibiting Bcl-2-like pro-survival proteins." *The Journal of Cell Biology* 186.3 (Aug. 2009), pp. 355–362.
- Mirschel, S., K. Steinmetz, M. Rempel, M. Ginkel, and E. D. Gilles (2009). "PROMOT: modular modeling for systems biology." *Bioinformatics (Oxford, England)* 25.5 (Mar. 2009), pp. 687–689.
- Moraru, I. I., J. C. Schaff, B. M. Slepchenko, and L. M. Loew (2002). "The virtual cell: an integrated modeling environment for experimental and computational cell biology." *Annals of the New York Academy of Sciences* 971 (Oct. 2002), pp. 595–596.
- Muzzey, D., C. A. Gómez-Urbe, J. T. Mettetal, and A. van Oudenaarden (2009). "A systems-level analysis of perfect adaptation in yeast osmoregulation." *Cell* 138.1 (July 2009), pp. 160–171.
- Neumann, L., C. Pforr, J. Beaudouin, A. Pappa, N. Fricker, P. H. Krammer, I. N. Lavrik, and R. Eils (2010). "Dynamics within the CD95 death-inducing signaling complex decide life and death of cells." *Molecular Systems Biology* 6.1, p. 352.
- O'Connor, C. L., S. Anguissola, H. J. Huber, H. Dussmann, J. H. M. Prehn, and M. Rehm (2008). "Intracellular signaling dynamics during apoptosis execution in the presence or absence of X-linked-inhibitor-of-apoptosis-protein." *Biochimica et biophysica acta* 1783.10 (Oct. 2008), pp. 1903–1913.
- Oliphant, T. E. (2007). "Python for Scientific Computing". *Computing in Science & Engineering* 9.3, pp. 10–20.
- Pedersen, M. and G. Plotkin (2008). "A language for biochemical systems". *Computational Methods in Systems Biology*, pp. 63–82.
- Rehm, M., H. J. Huber, C. T. Hellwig, S. Anguissola, H. Dussmann, and J. H. M. Prehn (2009). "Dynamics of outer mitochondrial membrane permeabilization during apoptosis". *Cell Death and Differentiation* 16.4 (Apr. 2009), pp. 613–623.

- Rehm, M., H. Dussmann, R. U. Janicke, J. M. Tavaré, D. Kogel, and J. H. M. Prehn (2002). "Single-cell fluorescence resonance energy transfer analysis demonstrates that caspase activation during apoptosis is a rapid process. Role of caspase-3". *The Journal of biological chemistry* 277.27 (July 2002), pp. 24506–24514.
- Rehm, M., H. J. Huber, H. Dussmann, and J. H. M. Prehn (2006). "Systems analysis of effector caspase activation and its control by X-linked inhibitor of apoptosis protein". *The EMBO Journal* 25.18 (Sept. 2006), pp. 4338–4349.
- Ren, D. et al. (2010). "BID, BIM, and PUMA are essential for activation of the BAX- and BAK-dependent cell death program". *Science (New York, NY)* 330.6009 (Dec. 2010), pp. 1390–1393.
- Schleich, K., U. Warnken, N. Fricker, S. Öztürk, P. Richter, K. Kammerer, M. Schnölzer, P. H. Krammer, and I. N. Lavrik (2012). "Stoichiometry of the CD95 death-inducing signaling complex: experimental and modeling evidence for a death effector domain chain model." *Molecular Cell* 47.2 (July 2012), pp. 306–319.
- Smith, A. M., W. Xu, Y. Sun, J. R. Faeder, and G. E. Marai (2012). "RuleBender: integrated modeling, simulation and visualization for rule-based intracellular biochemistry." *BMC Bioinformatics* 13 Suppl 8, S3.
- Sneddon, M. W., J. R. Faeder, and T. Emonet (2011). "Efficient modeling, simulation and coarse-graining of biological complexity with NFsim." *Nature Methods* 8.2 (Feb. 2011), pp. 177–183.
- Spencer, S. L., S. Gaudet, J. G. Albeck, J. M. Burke, and P. K. Sorger (2009). "Non-genetic origins of cell-to-cell variability in TRAIL-induced apoptosis." *Nature* 459.7245 (May 2009), pp. 428–432.
- SymPy Development Team (2012). *SymPy: Python library for symbolic mathematics*. Jan. 2012.
- Tan, C., P. J. Dlugosz, J. Peng, Z. Zhang, S. M. Lapolla, S. M. Plafker, D. W. Andrews, and J. Lin (2006). "Auto-activation of the apoptosis protein Bax increases mitochondrial membrane permeability and is inhibited by Bcl-2". *The Journal of biological chemistry* 281.21 (May 2006), pp. 14764–14775.
- Tiger, C.-F., F. Krause, G. Cedersund, R. Palmér, E. Klipp, S. Hohmann, H. Kitano, and M. Krantz (2012). "A framework for mapping, visualisation and automatic model creation of signal-transduction networks." *Molecular Systems Biology* 8.1, p. 578.
- Waltemath, D. et al. (2011). "Reproducible computational biology experiments with SED-ML—the Simulation Experiment Description Markup Language." *BMC systems biology* 5, p. 198.
- Willis, S. N., L. Chen, G. Dewson, A. Wei, E. Naik, J. I. Fletcher, J. M. Adams, and D. C. S. Huang (2005). "Proapoptotic Bak is sequestered by Mcl-1 and Bcl-xL, but not Bcl-2, until displaced by BH3-only proteins". *Genes & Development* 19.11 (June 2005), pp. 1294–1305.
- Willis, S. N. et al. (2007). "Apoptosis initiated when BH3 ligands engage multiple Bcl-2 homologs, not Bax or Bak". *Science (New York, NY)* 315.5813 (Feb. 2007), pp. 856–859.
- Wolfram Research, Inc. *Mathematica*. Version 8.0. Champaign, Illinois: Wolfram Research, Inc.

- Xu, T.-R. et al. (2010). “Inferring signaling pathway topologies from multiple perturbation measurements of specific biochemical species.” *Science signaling* 3,113, ra20.
- Youle, R. J. and A. Strasser (2008). “The BCL-2 protein family: opposing activities that mediate cell death”. *Nature Reviews Molecular Cell Biology* 9,1, pp. 47–59.

Programming biological models in Python using PySB

Supplementary Information

Carlos F. Lopez^{1,2,3}, Jeremy L. Muhlich^{1,2}, John A. Bachman^{1,2} and Peter K. Sorger^{2,4}

¹ Authors contributed equally to this work

² Center for Cell Decision Processes, Harvard Medical School, Department of Systems Biology, 200 Longwood Ave, Boston MA 02115, USA

³ Current address: Department of Cancer Biology, Center for Quantitative Sciences, Vanderbilt University School of Medicine, 2220 Pierce Avenue, Nashville TN 37232-6848, USA

⁴ To whom correspondence should be addressed

Email: Carlos Lopez - c.lopez@vanderbilt.edu; Jeremy Muhlich - jeremy_muhlich@hms.harvard.edu; John Bachman - bachman@fas.harvard.edu; Peter Sorger - peter_sorger@hms.harvard.edu;

Contents


Supplementary note	2
PySB syntax	4
Figure S1: Example macro implementations	11
Figure S2: Visualization of molecular species	12
Figure S3: UML class diagram	17
References	18

Supplementary note

In examining the differential equations for the models in Cui et al. [1], we found the following two errors.

Misplaced consumption term for activator

The reaction for the activation of Bax is described in Table 1 of [1] (“Chemical reaction network scheme”) as



In the list of ODEs provided in Table 2, the term describing the velocity of this reaction is listed (correctly) as

$$J_1 = k_1 \cdot [\text{InBax}] \cdot [\text{Act}]$$

The term J_1 appears in the following equations shown in Table 2 (some terms omitted for clarity):


$$\begin{aligned} d[\text{InBax}]/dt &= J_{\text{InBax}} - J_1 - J_9 \\ d[\text{AcBax}]/dt &= J_{\text{AcBax}} + J_1 - J_2 - \dots \\ d[\text{Act}]/dt &= J_{\text{Act}} - J_1 - J_3 + \dots \end{aligned}$$

The error is that J_1 appears as a negative term in the equation for $[\text{Act}]$, implying that activator is actually consumed in the reaction $\text{InBax} + \text{Act} \rightarrow \text{AcBax} + \text{Act}$. However, because this is a one-step “catalytic” reaction, with no intermediate complex formed, activator should not be consumed.

An additional indication that this is an error is that in the equations for the “Direct” model from Chen et al. [2], on which the Cui et al. [1] models are based, the equation for Activator does not have this negative term. The appearance of this error in the derived model, but not the original model, highlights the tendency of “copy-and-paste” model reuse to introduce inadvertent errors.

Missing term for Bax inactivation

The reaction for Bax inactivation is listed in Table 1 of Cui et al. [1] (“Chemical reaction network scheme”) as



In the list of ODEs provided in Table 2, the term describing the velocity of this reaction is J_5 :

$$J_5 = k_8 \cdot [\text{AcBax}]$$

J_5 appears correctly as a negative term in the equation for activated Bax:

$$d[\text{AcBax}]/dt = J_{\text{AcBax}} + J_1 - J_2 - J_4 - J_5 + J_8 - J_9 - 2 \cdot J_{10}$$

However, the equation for inactive Bax omits J_5 completely, where it should be incorporated as a positive term:

$$d[\text{InBax}]/dt = J_{\text{InBax}} - J_1 - J_9$$

This means that active Bax is consumed with the rate defined by J_5 , but the corresponding quantity of inactive Bax is not restored, leading to a loss of Bax over time. As with the error in the equation for Activator described above, the equation for Bax is correct in the original model from Chen et al. [2], indicating that the error was likely introduced in the process of duplicating the original model, or in the transcription of ODEs for publication.

PySB syntax

PySB model definition statements support some idiomatic syntax elements which may appear unusual to a Python programmer, but which provide economy of expression and a close mapping to the syntax of established rule-based modeling languages BNGL and Kappa. PySB defines overloaded operators and a so-called “SelfExporter” system which help give PySB model definition statements the feel of a domain-specific language like BNGL or Kappa within the general-purpose Python programming language. Modelers are not required to use the overloaded operators and SelfExporter functionality, but they are highly recommended for ease of readability.

Overloaded operators simplify rule expression syntax

Python’s built-in operators generally work with just a small set of built-in types, and an exception will be raised if they are applied to objects of any other class. For example “+” applies to numeric types (implementing mathematical addition), and separately to lists and strings (implementing concatenation). However a class may implement several special methods to *overload* these operators and allow them to be applied to its instances. A class is free to define any desired semantics in its overloaded operator implementations, but operator precedence and arity is fixed by the Python language grammar and cannot be altered via operator overloading.

PySB defines several classes representing nodes in an abstract syntax tree (AST) representation of a rule in the BNGL/Kappa domain-specific languages (DSL). These classes define overloaded operators allowing a modeler to write a Python expression which is visually very similar to one of these rules and evaluates to the corresponding AST. Whereas BNGL and Kappa use a standalone program to parse and simulate a model defined in a DSL, The PySB AST classes with operator overloading allow model components to be declared with Python program statements which evaluate directly to object representations of the given components. In this way Python itself serves as the parser for PySB models, and model components become live Python objects. The following table lists the AST classes and the syntactic elements they represent:

Class	Description
MonomerPattern	A pattern which matches instances of a given monomer, with optional restrictions on the state of certain sites.
ComplexPattern	A bound set of MonomerPatterns, i.e. a pattern to match a complex.
ReactionPattern	A pattern for the entire product or reactant side of a rule.
RulePattern	A container for the reactant and product patterns of a rule expression.

The operators to overload were chosen to fulfill two requirements. First, they must have the proper relative precedence in the Python grammar to minimize the need for parentheses and keep rule expressions uncluttered. Second, they should have a visual appearance as close as possible to the operators used in BNG and Kappa to help maintain consistency within the rule-based modeling ecosystem. The following table lists the overloaded operators and their semantics:

Operator	Description
()	Apply site conditions to a Monomer to create a MonomerPattern
%	Combine MonomerPatterns to create a ComplexPattern
+	Combine ComplexPatterns to create a ReactionPattern
<>	Combine two ReactionPatterns to create a reversible RulePattern
>>	Combine two ReactionPatterns to create an irreversible RulePattern

For those who are familiar with BioNetGen Language (BNGL), here are some actual expressions in both PySB and BNG syntax (Kappa is similar to BNGL):

PySB expression	BNGL equivalent
R(a=None)	R(a)
R(a=1) % R(b=1)	R(a!1).R(b!1)
R(a=None) + R(b=None)	R(a) + R(b)
R(a=None) + R(b=None) <> R(a=1) % R(b=1)	R(a) + R(b) <-> R(a!1).R(b!1)
R(a=None) + R(b=None) >> R(a=1) % R(b=1)	R(a) + R(b) -> R(a!1).R(b!1)

Below is a formal grammar for PySB rule expressions. Symbols corresponding to the Python AST node classes are shown in bold, using the actual class name. Symbols which are self-explanatory such as “site-name” and “string” are not expanded further.

$$\begin{aligned}
\langle \mathbf{MonomerPattern} \rangle &\rightarrow \langle monomer \rangle \text{ ‘(’ } \langle site-conditions \rangle \text{ ‘)’} \\
\langle site-conditions \rangle &\rightarrow \langle site-name \rangle \text{ ‘=’ } \langle condition \rangle \text{ ‘,’ } \langle site-conditions \rangle \\
&| \langle site-name \rangle \text{ ‘=’ } \langle condition \rangle \\
&| \phi \\
\langle condition \rangle &\rightarrow \langle string \rangle \\
&| \langle bond-number \rangle \\
&| \text{ ‘(’ } \langle bond-number-list \rangle \text{ ‘)’} \\
&| \text{ ‘(’ } \langle string \rangle \text{ ‘,’ } \langle bond-number \rangle \text{ ‘)’} \\
\langle bond-number-list \rangle &\rightarrow \langle bond-number \rangle \text{ ‘,’ } \langle bond-number \rangle \\
&| \langle bond-number-list \rangle \text{ ‘,’ } \langle bond-number \rangle \\
\langle \mathbf{ComplexPattern} \rangle &\rightarrow \langle \mathbf{ComplexPattern} \rangle \text{ ‘%’ } \langle \mathbf{ComplexPattern} \rangle \\
&| \langle \mathbf{MonomerPattern} \rangle \\
\langle \mathbf{ReactionPattern} \rangle &\rightarrow \langle \mathbf{ReactionPattern} \rangle \text{ ‘+’ } \langle \mathbf{ReactionPattern} \rangle \\
&| \langle \mathbf{ComplexPattern} \rangle \\
\langle \mathbf{RulePattern} \rangle &\rightarrow \langle \mathbf{ReactionPattern} \rangle \langle rule-op \rangle \langle \mathbf{ReactionPattern} \rangle \\
\langle rule-op \rangle &\rightarrow \text{ ‘<>’ } | \text{ ‘>>’ }
\end{aligned}$$

Excessive or haphazard use of operator overloading can certainly lead to confusing code, but we felt the construction of rule ASTs was a reasonable application with a limited scope. For comparison, here are several PySB rule expressions written using both overloaded operators and explicit AST assembly. The explicit forms of the first four subexpressions look simple enough in isolation, but the economy of the overloaded operators becomes readily apparent upon considering the final full RuleExpression.

Operators	Explicit AST assembly
<code>mp = R(a=1)</code>	<code>mp = MonomerPattern(R, {'a': 1})</code>
<code>cp = mp1 % mp2</code>	<code>cp = ComplexPattern([mp1, mp2])</code>
<code>rp = cp1 + cp2</code>	<code>rp = ReactionPattern([cp1, cp2])</code>
<code>re = rp1 <> rp2</code>	<code>re = RuleExpression(rp1, rp2, True)</code>
<code>R(a=None) + R(a=None) <> R(a=1) % R(a=1)</code>	<pre> RuleExpression(ReactionPattern([ComplexPattern([MonomerPattern(R, {'a': None})]), ComplexPattern([MonomerPattern(R, {'a': None})])]), ReactionPattern([ComplexPattern([MonomerPattern(R, {'a': 1}), MonomerPattern(R, {'a': 1})])]), True) </pre>

SelfExporter functionality streamlines model construction

PySB also includes functionality to streamline the process of creating components and adding them to models, using a class called **SelfExporter**. Like all object constructors in Python, each of the component constructors (**Monomer**, **Rule**, **Parameter** and **Compartment**) return an instance of the requested component. In a typical programming paradigm, it would be necessary to explicitly retain a reference to the created object in a variable for later use. For example, creating a monomer “R” and parameter “kf” for use in a rule declaration would require the following statements:

```

R = Monomer('R', ['a'])
kf = Parameter('kf', 1)
dimerize = Rule('dimerize', R(a=None) + R(a=None) >> R(a=1) % R(a=1), kf)

```

Here the **Monomer** constructor is used to create an instance of a **Monomer** object named “R”, stored in the local variable **R**. From a modeling perspective, one can immediately see a potentially confusing aspect of this approach: we now have to mentally keep track of two “names” for the same monomer, one the variable storing the reference to the object (**R**, which must be used to

build up the expression for the dimerization rule) and one the descriptive name assigned to the new object (“R”). Even though we have chosen to use the same name for both the object’s descriptive name and its variable in order to minimize confusion, maintaining this consistency requires mental effort on the part of the modeler and clutters the code making it harder to read.

In addition to managing the issue of naming, we must also add the newly created **Monomer**, **Parameter** and **Rule** objects to a model. To do this, we must call the `model.add_component` method on each component object:

```
model = Model('model')
model.add_component(R)
model.add_component(kf)
model.add_component(dimerize)
```

This repetition adds further visual noise to the model code, and accidentally omitting the `add_component` call for one or several components can lead to errors far from the site of declaration (in the case of a **Monomer** or **Parameter** used in a distant **Rule**) or worse, subtle errors in model behavior (in the case of a **Rule**).

In a typical modeling scenario, creation of model components tends to follow the pattern described above, that is:

1. Create a component using the appropriate constructor and assign it to a variable in the current namespace.
2. Add the created component to the current model.

The repetition of this pattern for every component in a model tends to be verbose and obscure the model structure; it also creates opportunities for error as described above.

PySB includes a helper class called **SelfExporter** that streamlines model definition by automatically performing the above steps. Using the functionality provided behind the scenes by **SelfExporter**, we can now simply write:

```
Model('model')
Monomer('R', ['a'])
Parameter('kf', 1)
Rule('dimerize', R(a=None) + R(a=None) >> R(a=1) % R(a=1), kf)
```

In the above example, when the constructor `Model('model')` is called, the **SelfExporter** “exports” a reference to the model by *creating a global variable called `model` in the current namespace and assigning a reference to the created **Model** object to it*. (This is possible because by using the Python module **inspect**, global variables in any accessible namespace can be accessed and manipulated programmatically: they are stored as a dictionary linking the name of the variable (“model”) to its value (a reference to the new **Model** object). The **SelfExporter** can add global variables by modifying the entries in this dictionary.)

When the second statement, `Monomer('R', ['a'])`, is executed, the `SelfExporter` performs a similar action: it creates a new variable, `R`, and assigns to it the reference to the new object (a `Monomer` object given the name “R”). However, in this case it also takes a second action: it adds the new `Monomer` object `R` to the set of `Monomer` objects associated with the currently defined model, `model`. The process for the `Parameter` object is exactly the same: a global variable `kf` for the new object is created and added to `model`.

Finally, because they have been “exported” as variables by the behind-the-scenes action of the `SelfExporter`, the `Monomer` variable `R` and the `Parameter` variable `kf` are now globally accessible, and we can use both in the `Rule(...)` definition that follows. The `Rule` object itself is similarly exported and added to the model.

To summarize, the execution of the above code results in:

- The addition of four variables to the global namespace: `model`, `R`, `kf`, and `dimerize`
- The addition of the model components `R`, `kf`, and `dimerize` to the model `model`.

It should be noted that though it is the default behavior, *the use of `SelfExporter` functionality is entirely optional*. In certain sophisticated modeling scenarios involving the dynamic creation of multiple alternative models, the explicit approach to component creation and assignment may be preferred. However, we have found that the `SelfExporter` substantially simplifies the most common modeling use cases. A summary of the syntax for the simple example described above, with and without the action of `SelfExporter`, is shown below:

With <code>SelfExporter</code>	Without <code>SelfExporter</code>
<code>Model('model')</code>	<code>model = Model('model')</code>
<code>Monomer('R', ['a'])</code>	<code>R = Monomer('R', ['a'])</code> <code>model.add_component(R)</code>
<code>Parameter('kf', 1)</code>	<code>kf = Parameter('kf', 1)</code> <code>model.add_component(kf)</code>
<code>Rule('dimerize', R(a=None) + R(a=None) >> R(a=1) % R(a=1), kf)</code>	<code>dimerize = Rule('dimerize', R(a=None) + R(a=None) >> R(a=1) % R(a=1), kf)</code> <code>model.add_component(dimerize)</code>

The following comparison with the BNGL syntax for the same simple model shows how the combined use of overloaded operators and the `SelfExporter` give models written in PySB the feel of a domain-specific language embedded in Python:

PySB statement	BNGL equivalent
Model('model')	<i>(not needed)</i>
Monomer('R', ['a'])	begin molecule types R(a) end molecule types
Parameter('kf', 1)	begin parameters kf 1 end parameters
Rule('dimerize', R(a=None) + R(a=None) >> R(a=1) % R(a=1), kf)	begin reaction rules dimerize: R(a) + R(a) -> R(a!1).R(a!1) kf end reaction rules

Figure S1

A) catalyze macro basic implementation


Basic Implementation	<pre>def catalyze(enz, e_site, sub, s_site, prod, klist): kf, kr, kc = klist # Get the parameters from the list # Create the rules rb = Rule('bind_%s_%s' % (enz().monomer.name, sub().monomer.name), enz({e_site:None}) + sub({s_site:None}) <> enz({e_site:1}) % sub({s_site:1}), kf, kr) rc = Rule('catalyze_%s%s_to_%s' % (enz().monomer.name, sub().monomer.name, prod().monomer.name), enz({e_site:1}) % sub({s_site:1}) >> enz({e_site:None}) + prod({s_site:None}), kc) return [rb, rc]</pre>
----------------------	--

B) catalyze_one_step macro


Example Macro Call	<pre>catalyze_one_step(C8(bf=None), Bid(state='U', bf=None), Bid(state='T', bf=None), kf)</pre>
Basic Implementation	<pre>def catalyze_one_step(enz, sub, prod, kf): # Create the rule r = Rule('catalyze_one_step_%s_%s_to_%s' % (enz.monomer.name, sub.monomer.name, prod.monomer.name), enz() + sub() >> enz() + prod(), kf) return r</pre>
BNGL Rules	<pre>C8(bf) + Bid(bf,state~U) -> C8(bf) + Bid(bf,state~T) kf</pre>
ODEs	<pre>C8: ds0/dt = 0 Bid: ds1/dt = -kf*s0*s1 tBid: ds2/dt = kf*s0*s1</pre>

(A) Simplified implementation of the `catalyze` macro. The `Rule` objects for the binding and catalytic steps are created according to defined templates, with the species identities (enzyme, substrate, and product), binding site names, and parameters filled in from the arguments to the macro. (B) `catalyze_one_step`. This macro models a “one-step” approximation of catalysis according to the reaction scheme $E + S \rightarrow E + P$. The macro creates a single catalysis rule according to the prescribed template, which can then be used to generate the BNGL rule and set of ODEs shown below. The full implementation of the `catalyze` and `catalyze_one_step` macros, with documentation and handling of various special cases, can be found in the `macros.py` file in the PySB source code online (<http://pysb.org>).


Figure S2




continued...




continued...



continued...




continued...



Output from the PySB `render_species` tool run against EARM 2.0-M1a (Lopez Embedded). Each large box represents one species, with its number and PySB representation at the top, followed by a depiction of the monomer graph. In the monomer graph, each segmented box represents a monomer, with its name in the first green segment and its sites in the following segments. Edges between monomer sites represent bonds, and their numeric labels correspond directly to the bond numbering in the PySB representation above.

Figure S3



UML class diagram for core PySB classes. Note that for clarity, only fields are shown (not methods). Model objects are built up using collections of Rule, Parameter, Compartment, Observable, and Monomer objects, which acquire generic functionality (such as a name field) from the superclass Component. Components are aggregated in the collection class ComponentSet. Rules are expressed as instances of the class RuleExpression, which are built up using instances of the classes ReactionPattern, ComplexPattern, and MonomerPattern. The resulting RuleExpression is stored in a Rule object, along with the associated rate parameters (as Parameter objects). Additional information about model elements (including Components or macros) can be stored in an Annotation object. The list of Annotations is stored in the Model.

References

1. J. Cui, C. Chen, H. Lu, T. Sun, P. Shen, *PLoS ONE* **3**, e1469 (2008).
2. C. Chen, J. Cui, W. Zhang, P. Shen, *FEBS letters* **581**, 5143 (2007).

What makes Michaelis and Menten’s model so significant was not that it fits the experimental data, but that it provides evidence for something unseen.

Jeremy Gunawardena (Gunawardena, 2012)

3

Cancer-associated Bax point mutations block apoptotic pore formation at distinct conformations

3.1	Introduction	87
3.2	Results	89
3.3	Discussion	104
3.4	Methods	106
3.5	Supplementary information	110

AUTHORSHIP CONTRIBUTION

I performed analysis and modeling of the kinetic data; experiments were performed by J. Kale. I wrote this draft of the manuscript following an outline developed in collaboration with J. Kale.

ABSTRACT

Mitochondrial outer membrane permeabilization is a key checkpoint in apoptosis determining whether cells live or die, but the dynamics and ordering of the steps leading to pore formation are not well characterized. In this study we use kinetic analysis of a reconstituted protein-membrane system to characterize the dynamics of the structural rearrangements of the pore-forming effector Bax as it undergoes activation by either cBid or Bim, followed by membrane insertion and pore formation. Using site-specific labeling of different regions of the protein with the environmentally sensitive dye NBD, we find that the protein undergoes a rearrangement through at least one structurally distinct intermediate en route to the final inserted state. This intermediate involves transient increases in the hydrophobicity of residues in the BH₃ region and is a necessary precursor to pore formation. FRET measurements indicate that the initial conformational change of Bax corresponds closely to the formation of the Bax:BH₃-only activator complex, while the second transition corresponds to the formation of Bax oligomers. While this sequence of rearrangements for Bax was the same for activation by both cBid and Bim, cBid appeared to undergo a second rearrangement relative to Bax whereas Bim did not. To determine whether the pore-forming intermediate we identified has functional significance, we prepared mutants of Bax identified from cancer genome sequencing and characterized them in our *in vitro* system. We found that three cancer-associated mutants appeared to be trapped in the intermediate state with diminished pore formation, suggesting that transitions into and out of this intermediate can be modulated independently with consequences for the execution of apoptosis. The accompanying dataset, involving corresponding dye release, Bax insertion, and FRET timecourses between Bax and activators Bid and Bim, is provided in a variety of formats for reuse and re-analysis.

3.1 INTRODUCTION

Apoptosis is a biochemically programmed cell death pathway that governs when and how cells die in response to internal or external stress stimuli. A key step in this process is mitochondrial outer membrane permeabilization, or MOMP, in which pores form in the mitochondrial outer membrane and release pro-apoptotic proteins, such as cytochrome *c* and Smac, into the cytosol (Tait and Green, 2010). These proteins activate a family of proteases known as caspases, which trigger the rapid degradation of cellular components. The processes downstream of MOMP are in general rapid and complete, making MOMP itself a key regulator of cell fate (Albeck et al., 2008; Goldstein et al., 2000; Rehm et al., 2002). Functional studies have shown that the tendency of mitochondria derived from cell lines and tumor cells to undergo MOMP correlates strongly with cell death in response to cytotoxic chemotherapy and ultimately, clinical outcomes (Chonghaile et al., 2011; Vo et al., 2012).

MOMP is regulated by a set of evolutionarily related proteins known as the Bcl-2 family, which are responsible for activating, inhibiting, and forming the pores in the mitochondrial outer membrane (Chipuk et al., 2010). Anti-apoptotic family members include the known oncogenes Mcl-1, Bcl-XL, and Bcl-2; pro-apoptotic members are divided into the “effectors” Bax and Bak, which form the pores in the mitochondrial membrane, and the larger group of “BH₃-only proteins,” which are so called because they share only the third of the four Bcl-2 homology (BH) regions with the other family members. The BH₃-only proteins promote apoptosis by both inhibiting the anti-apoptotic family members and activating the pore-forming effectors, though with differential affinity and specificity (Letai et al., 2002; Certo et al., 2006). A number of chemical modulators of the Bcl-2 family members have been developed, including inhibitors of the anti-apoptotic members (Oltersdorf et al., 2005; Souers et al., 2013; Levenson et al., 2015) and agonists of the pro-apoptotic effectors (Gavathiotis et al., 2012).

Much is now known about the structural basis of Bcl-2 protein interactions, and this knowledge has proven essential to the efforts to develop targeted therapeutics. Prior to activation, a fraction of Bax binds transiently to membranes but it is predominantly cytosolic (Yethon et al., 2003). In this soluble conformation, the hydrophobic groove of Bax is occluded by the C-terminal α_9 helix (Suzuki et al., 2000). Engagement of a BH₃-only protein such as Bim at a “trigger site” on the rear surface of the protein leads to

displacement of the C-terminal helix, exposure of the 6A7 epitope, and membrane insertion (Kim et al., 2009; Gavathiotis et al., 2010). Other studies have shown that binding of the BH₃-only protein to the hydrophobic groove is a necessary step for activation and insertion of Bax (Dai et al., 2011; Czabotar et al., 2013).

Several recent structural studies have given a clearer picture of the arrangement of Bax and Bak molecules at the assembled pore. The dimerization interface of Bax has been mapped by mutagenesis (George et al., 2007), and recent studies have suggested that Bax dimerization involves symmetric binding via the BH₃ region and hydrophobic groove in a fashion similar to how these proteins bind with anti-apoptotics (Czabotar et al., 2013; Bleicken et al., 2014).

Despite this new information, a number of important questions regarding apoptotic pore formation remain. Existing studies based on static or equilibrium measurements offer insight into the structure of the assembled pore but less information about the nature of any intermediate conformations taken by Bax or Bak. Conversely, kinetic studies have clarified aspects of the sequence of steps taken en route to pore formation, but lack structural information about the conformational states of Bax during this process (Lovell et al., 2008; Kushnareva et al., 2012; Saito et al., 2000). What is needed is a dynamic picture of the process by which Bax goes from the aqueous state to the assembled pore, using full-length proteins and membranes.

An additional question pertains to the interaction between Bax and its BH₃-only activators during activation and pore formation. While it has been shown that tBid remains bound to Bax at equilibrium, the significance of this post-activation complex is unknown (Lovell et al., 2008; Kim et al., 2009). In addition, because several of the existing structural studies have used a variety of methods for activating Bax (e.g., BH₃ peptides vs. full-length BH₃-only proteins and liposomes vs. CHAPS detergent), it is unclear whether the particular rearrangements of Bax are contingent on the activator used.

In this study we use kinetic analysis of a reconstituted system to characterize the conformational changes of Bax en route to pore formation. We accomplish this by labeling the protein at 19 positions with an environmentally sensitive dye that reports on membrane insertion, complemented by intermolecular FRET to determine the timing of protein-protein interactions. We analyze the approximately 400 experimental

timecourses using a set of kinetic models that allows us to identify the number and timing of conformational changes. We find that Bax transitions rapidly to an intermediate associated with binding of cBid or Bim, followed by a slower transition associated with Bax dimerization and pore formation. Mutations found in cancer cells can trap Bax in this intermediate, preventing MOMP. The accompanying dataset, involving corresponding dye release, Bax insertion, and FRET timecourses between Bax and activators cBid and Bim, is provided in a variety of formats along with source code at <https://github.com/johnbachman/tbidbaxlipo>. Plots of raw and processed data can be accessed in the Supplementary Online Material (SOM) at <http://sorger.med.harvard.edu/data/bachman/kale/index.html>.

3.2 RESULTS

To characterize the structural rearrangements associated with Bax insertion and oligomerization, we generated a set of Bax mutants labeled with the environment-sensitive dye NBD. Because NBD fluorescence increases in a hydrophobic environment, the NBD-labeled Bax mutants can be used to track the rate and extent of membrane insertion and other structural rearrangements kinetically (Lovell et al., 2008). Site-specific NBD labeling of the protein was accomplished by mutating the two endogenous cysteines of Bax to serine and introducing single cysteine mutations at other positions (Methods). In total, 19 Bax mutants were purified and labeled with NBD, with the labeled positions spanning the N-terminus, BH₃ region, pore-forming helices (α_5 and α_6), and C-terminus (α_9) (Figure 3.1A-B). NBD labeling efficiency ranged from 56% to 92%, with most mutants over 75% (SOM Section 2).

We incubated the mutants with synthetic liposomes and a BH₃-only activator (cBid or Bim) and tracked the kinetics of NBD-Bax fluorescence and liposome permeabilization simultaneously using a multiplexed fluorescence spectroscopy approach described previously (Figure 3.1C) (Lovell et al., 2008; Kale et al., 2014). For each condition, three replicate timecourses were collected, each using liposomes from separate preparations. The initial kinetic dataset consisted of 20 Bax alleles \times 2 activators \times 3 replicates = 120 distinct kinetic experiments, with two fluorescence channels for each of the 19 mutants (NBD and terbium), for a total of 234 kinetic curves.

As measured by dye release at the 1 hr endpoint, two thirds of the mutants had approximately wild-




Figure 3.1: Assaying dye release and conformational changes of labeled Bax mutants.

(A) Primary and secondary structure of Bax.

(B) Solution structure of Bax (PDB ID: 1F16) with positions of engineered cysteine mutations used for dye labeling.

(C) Design of multiplexed fluorescence spectroscopy assays. A BH3-only activator (which in some experiments is labeled with the fluorescence donor DAC) is added to a solution of synthetic lipid vesicles that contain the fluorescent ion terbium and allowed to equilibrate. NBD-labeled Bax mutants are then added and fluorescence timecourses are recorded. Pore formation is measured as a decrease in terbium fluorescence as it is released into the solution, which contains its quencher EDTA; Bax conformational changes and membrane interactions are indicated by the fluorescence changes of NBD-Bax; and interactions between Bax and its activators are given by the quenching of the DAC-labeled BH3-only protein by NBD-Bax.

(D) Endpoint dye release activity of the NBD-labeled Bax mutants activated by either unlabeled cBid or Bim, normalized to the activity of wild-type Bax.

(E) Endpoint NBD-Bax fluorescence intensities after activation by either unlabeled cBid or Bim, relative to the initial intensity F_i .

type levels of activity (80% or more); a subset of mutants was slightly less active (SOM Section 1.1). The kinetic curves can be inspected in the plots in SOM Section 3.1.

3.2.1 NBD-LABELED BAX MUTANTS UNDERGO DISTINCT CHANGES IN HYDROPHOBICITY INDICATING WIDESPREAD CONFORMATIONAL CHANGE

The labeled mutants exhibited a variety of NBD fluorescence intensities at the timecourse endpoint, indicative of the membrane-bound conformation of Bax in assembled pores (Figure 3.1D). Consistent with prior studies, the $\alpha 5$ residues 120, 122, and 126 as well as the C-terminal residues 175 and 179 had high NBD fluorescence values, suggestive of membrane insertion (Annis et al., 2005; Westphal et al., 2014). However, since the NBD fluorescence is influenced not only by interactions with membranes but also by the side chains of neighboring amino acids and other intramolecular effects, it is an incomplete measure of the degree of conformational change. The relative NBD fluorescence values, which indicate the fold-change in NBD fluorescence at the labeled residue relative to the initial aqueous state, are shown in Figure 3.1E; a value greater than one indicates an increase in fluorescence, and thus hydrophobicity; a value less than one, a decrease. All labeled residues showed a relative increase in NBD fluorescence at equilibrium, with the exception of the mutant labeled at position 47, which decreased. A number of residues, including 15, 62, 184, and 188, had relatively low absolute NBD fluorescence but large relative fluorescence changes, indicating that the rearrangements of Bax during pore formation encompass most regions of the protein, not only those in the pore-forming and C-terminal helices.

3.2.2 THE KINETICS OF NBD FLUORESCENCE CHANGES FOR BH₃ RESIDUES SUGGEST AN INTERMEDIATE CONFORMATIONAL STATE

To investigate the dynamic rearrangements of Bax occurring during activation and pore formation we examined the kinetics of NBD fluorescence changes. Our expectation, based on a previous study using NBD-126C-Bax, was that each mutant would exhibit single-exponential kinetics that would allow each residue to be ranked in order of its insertion timescale, thereby describing an insertion “pathway” (Lovell et al., 2008).

The kinetic data showed that the mutants had distinct rates and dynamics of NBD fluorescence changes,




Figure 3.2: Kinetics of NBD fluorescence changes.

(A) Example NBD-Bax fluorescence timecourses, normalized to the initial fluorescence F_0 . Each of the six timecourses shown used cBid as the BH3-only activator and is the first of three replicates.

(B) Percentage of maximum NBD fluorescence reached at an early timepoint (105 sec), plotted against the amount of dye release relative to wild-type Bax, using cBid as the activator.

(C) As for (B), but using Bim as the activator.

(D) Derivatives of NBD and terbium release timecourses for NBD-15C-Bax (first replicate, cBid as activator), normalized to the maximum rate of change. Derivatives were calculated numerically after processing the curves with a low-pass filter to reduce noise. The curves shown were for the first experimental replicate with cBid as the activator.

(E) As for (D), but for NBD-54C-Bax (first replicate, cBid as activator). The vertical grey line is drawn at the peak rate of change of terbium release.

falling into two classes depending on the location of the label. While the majority of the residues showed monotonic increases in fluorescence similar to what we had observed previously, three residues near the BH3 region (54, 68, 79) showed a transient peak in fluorescence followed by a decline towards the equilibrium value (Figure 3.2A). The rates of fluorescence change varied among the mutants with monotonic kinetics, with some (e.g., NBD-3C-Bax) approaching their equilibrium value rapidly and others (e.g., NBD-179C-Bax) much more slowly (Figure 3.2A). As a measure of the differences in NBD kinetics across all mutants we calculated the extent to which the relative NBD fluorescence had attained its maximal value at an early timepoint (105 sec) and found it to vary considerably among the labeled residues (SOM Section 1.1).

A corresponding analysis of dye release levels at the same early timepoint showed that the mutants were also highly variable in their overall rates of pore formation, preventing a straightforward comparison of NBD insertion timescales (SOM Section 1.1). Most Bax mutants, including those that had endpoint release activity similar to wild-type, showed accelerated dye release kinetics, suggesting that for many mutants the addition of the NBD label preserved Bax's pore formation activity but destabilized the protein by reducing energy barriers between its conformational states.

To assess the contribution of activity differences to the rates of NBD fluorescence change, we plotted the NBD kinetics data against the dye release data for the 105 sec timepoint and found that the residues fell into a set of two classes (Figures 3.2B-C). For residues outside the BH₃ region, faster rates of NBD fluorescence change corresponded to faster rates of dye release in a roughly linear pattern, suggesting that most of the differences in NBD fluorescence at 105 seconds could be explained via differences in dye release activity resulting from the NBD label. However, residues in or near the BH₃ region, including 54, 62, 68, and 79, showed significant early NBD fluorescence changes that differed considerably from what would be expected due to activity differences alone. This suggests that the rapid increase in the fluorescence of these labeled residues is due to a structural rearrangement that precedes dye release.

To further analyze the timing of NBD fluorescence changes relative to dye release, we calculated the derivatives of both curves (SOM Section 3.3). While the rate of change of NBD fluorescence is almost invariably fastest at the initial timepoints, the rate of change of dye release reaches a maximum before declining to zero at equilibrium (shown for NBD-15C-Bax in Figure 3.2D; other residues in SOM Section 3.3). This process is evident in the original dye release timecourses as a slight lag phase at the earliest timepoints that is also present for wild-type Bax (SOM Section 3.1). The appearance of the lag phase in dye release was consistent for all mutants except a handful with low endpoint activity (68, 79, 179, and 188), though the duration of the lag between the labeled mutants varied depending on their release kinetics (mutants with accelerated release kinetics had a shorter lag phase). The analysis of the timecourse derivatives indicates that even for the residues outside of the BH₃ region with monotone kinetics, NBD fluorescence at the labeled residues undergoes an initial change that precedes pore formation.


Strikingly, for NBD-54C-Bax, a BH₃-region mutant exhibiting both non-monotone kinetics and a

measurable lag phase in dye release, the timepoint at which the maximum NBD fluorescence is reached (the zero of the NBD derivative) corresponds very closely to the timepoint at which dye release achieves its maximum rate (the peak in the dye release derivative; vertical line, [Figure 3.2E](#)). One possible interpretation is that the rate of dye release is proportional to the abundance of an intermediate species associated with increased hydrophobicity at the Bax BH₃ region. However, a more rigorous quantitative analysis is required to determine whether this relationship between NBD fluorescence and pore formation is broadly consistent with the data.

Taken together, the kinetic data indicate that many regions of Bax, particularly the BH₃ region, undergo a concerted transition to an intermediate state of increased hydrophobicity that precedes pore formation and dye release.

3.2.3 KINETIC ANALYSIS OF NBD FLUORESCENCE TIMECOURSES SUGGESTS THREE OR FOUR DISTINCT FLUORESCENCE STATES

Our analysis suggests the presence of at least one intermediate conformational state for Bax prior to pore formation. We sought to determine whether our data indicated additional fluorescence transitions, and the degree to which each of the labeled residues participated in these transitions. We therefore considered a simple kinetic model involving linked equilibria between conformational states of Bax, each with differential NBD fluorescence:




with NBD fluorescence defined by

$$NBD(t) = \sum_{i=1}^n C_i [Bax_{Fi}(t)] \quad (3.2)$$

where the C_i are scaling parameters denoting the NBD fluorescence of the various states of Bax. Though the generalized scheme in [Eq. 3.1](#) incorporates both the forward and reverse reactions, in our analysis we

considered models containing only irreversible forward reactions:



For experimental systems such as ours where only the aggregate fluorescence is measured, the reversible and irreversible formulations produce equivalent fits to the data and are indistinguishable in terms of their ability to identify the number of conformational states ([Supplementary information](#)).

Qualitative inspection suggests that the kinetic curves for several of the NBD-labeled residues are reasonably well approximated by the two-state model, which takes the form of a single-exponential equation ([Methods](#), [Eq. 3.8](#); [Supplementary information](#)). This is particularly true for residues outside the BH₃ region (e.g., residues 3, 120, and 179 shown in [Figure 3.3A](#); others in [SOM Section 3.4](#)). However, the non-monotone fluorescence curves of some BH₃-region residues (54, 68, 79) are very poorly fit by the single-exponential equation [Eq. 3.8](#) ([Figure 3.3A](#)). Indeed, they cannot be fit by any two-conformation scheme in the form of [Eq. 3.1](#) where the labeled residue undergoes a single fluorescence transition, due to the form of the mathematical formula describing such systems ([Supplementary information](#)).

We therefore defined an expanded set of kinetic models according to the reaction scheme in [Eq. 3.3](#) with either two, three, four, or five conformational states and conditioned them on the NBD fluorescence data using Bayesian parameter estimation ([Methods](#)). An advantage of the Bayesian approach is that given a suitable error model, the data can be used to rigorously estimate the posterior probability distributions of the model parameters, and can also discriminate between models with different numbers of parameters ([Gelman et al., 2014](#); [Girolami et al., 2010](#)). Prior probability distributions for the fluorescence scaling parameters were drawn from the observed relative and absolute dynamic ranges for the dye NBD; prior distributions for the kinetic parameters were uniform and bounded by the relevant timescales of the experiment (ranging from 10⁻¹ to 10⁻⁶ sec⁻¹).

Notably, after calibrating the models we found that all residues, not only those in the BH₃ region, showed evidence of at least three conformational states, shown by the dramatic improvement in both the fits ([Figure 3.3B](#)) and marginal likelihood values in going from two to three conformations ([Figure 3.3C](#)).




Figure 3.3: Fits of kinetic models to NBD fluorescence timecourses.

Figure 3.3: (Previous page.) Fits of kinetic models to NBD fluorescence timecourses.

(A) Fits of the two-conformation model to the curves in Figure 2A.

(B) Fits of the three-conformation model to the curves in Figure 2A.

(C) Marginal likelihood values calculated for different models for all NBD timecourses in the dataset by Markov chain Monte Carlo sampling. Inset shows differences between the three, four, and five conformation models on an increased scale. Boxes extend from the first to the third quartile of the data, with a horizontal line at the median; whiskers extend above and below to $1.5 \times$ the interquartile range.

(D) Marginal posterior distributions of the parameters of the three-conformation model. The color intensities along each horizontal red and green band indicate the MCMC sample frequency at the corresponding interval of the parameter distribution. In the C_2/C_3 plot, frequencies for C_2 are indicated in red, C_3 in green; in the k_1/k_2 plot, frequencies for k_1 are indicated in red, k_2 in green. The section of the plot for a given residue contains three corresponding pairs of red and green bands; each pair is drawn from fits to a single experimental replicate. Residues with identifiable parameter values are marked "id."

(E) Covariances of the marginal posterior distributions for NBD-54C-Bax + cBid, replicate 1, indicating the two different modes of fitting the data.

(F) Marginal posterior distributions of the parameters of the three-conformation model using a log-normal prior distribution for the scaling parameters C_2 and C_3 .

A handful of residues showed evidence of a fourth conformational state (Figure 3.3C and SOM Section 3.5). Notably, none of the residues showed evidence supporting five or more conformation (i.e., four or more fluorescence transitions).

3.2.4 FITTED MODEL PARAMETERS SUGGEST A FAST TRANSITION TO AN INTERMEDIATE CONFORMATION WITH RELATIVELY HIGHER HYDROPHOBICITY FOR THE BH₃ AND SURROUNDING REGIONS

To gain insight into the dynamics of conformational change across different regions of Bax, we inspected the posterior distributions of the parameters of the irreversible three-conformation model. The parameters of this model included the rates associated with the first and second transitions, k_1 and k_2 , respectively, as well as the fitted NBD fluorescence intensities of the intermediate and final conformations, C_2 and C_3 . The posterior distributions for these four parameters are summarized for all mutants and replicates with cBid as the activator in Figure 3.3D.

To determine which of the two transitions is faster under these experimental conditions, the first or the second, we examined the posterior distributions for the rate parameters k_1 and k_2 (Figure 3.3D). For nine of the mutants, we found that both parameters were clearly identifiable, indicated by tight, unimodal posterior distributions. For these mutants the ranking of k_1 and k_2 was consistent, with the first transition found to be at least ten-fold faster than the second transition.

The remaining mutants had considerable bimodality in their posterior distributions, admitting two types of fits: fast k_1 and slow k_2 , consistent with the identifiable set of mutants, or the converse, slow k_1 and fast k_2 . Examination of the covariance structure of the parameters showed that this latter mode of fitting was invariably associated with significantly greater predicted NBD fluorescence intensities for the intermediate conformation, C_2 (shown for NBD-54C-Bax in [Figure 3.3E](#); covariance plots for other residues in [SOM Section 3.4](#)).

In our estimation procedure, we had used a uniform prior distribution for the fluorescence parameters C_2 and C_3 , which assumed that *a priori*, values for these parameters falling within a 10-fold change in either direction from the initial value were equally likely. Given that the endpoint relative fluorescence values for most mutants was less than 2-fold, we revisited this assumption, recalibrating the three-conformation models using a more restrictive prior distribution for C_2 and C_3 . Using a lognormal distribution centered around the initial fluorescence value with a standard deviation of 1.5-fold change in either direction, we found that bimodality in the parameter posterior distributions was reduced, with posterior distributions consistently indicating values for k_1 substantially faster than k_2 ([Figure 3.3F](#)).

Using the more restrictive prior, we also analyzed the fitted fluorescence values C_2 and C_3 across the residues. With the exception of position 62, the residues in or surrounding the BH₃ region (positions 36-79) were associated with greater hydrophobicity in the intermediate relative to the final conformations, while the pore forming and C-terminal regions (positions 120-188) were associated with greater or equal hydrophobicity in the final conformation; the three residues in the N-terminus (3, 5, and 15) were mixed or ambiguous in their behavior.

3.2.5 BOTH cBID AND BIM REMAIN ASSOCIATED WITH BAX AT EQUILIBRIUM, BUT THE cBID:BAX INTERACTION PEAKS EARLY

To determine the timing of the conformational changes of Bax relative to its interactions with its BH₃-only activators, we selected the NBD-Bax mutants with near-wild-type levels of activity and measured their interactions with cBid and Bim by Förster resonance energy transfer (FRET). We labeled single-cysteine mutants of cBid and Bim with the donor fluorophore DAC at the h_o position of their respec-

tive BH₃ regions (Czabotar et al., 2013), which is quenched by the NBD acceptor upon binding with Bax (Lovell et al., 2008). The resulting FRET traces give a measure of the BH₃-only/Bax interaction timescales as well as the degree to which different residues of Bax remain proximal to Bid during the process of pore formation. Kinetic assays were performed for wild-type Bax and 12 mutants \times 2 activators \times 3 replicates = 72 assays with three channels for the labeled mutants, yielding 210 curves (plots of raw data in SOM Section 5.1). For several experimental replicates a number of outlier timepoints were observed in the FRET data, likely due to fluorescent debris in the solution; these outliers were removed manually in a preprocessing step (SOM Section 5.2).

As shown in Figure 3.4A, changes in FRET occurred on a relatively fast timescale, qualitatively similar to the initial changes in NBD fluorescence. Endpoint FRET levels between the BH₃-only proteins and Bax were variable for the different sites on Bax, with high levels at residue 36 and the α ₅ helix residues 122 and 126 for both DAC-cBid and DAC-Bim (Figure 3.4B). Strikingly, for several labeled Bax residues including NBD-54C, the DAC-cBid FRET reached an early peak and then declined, whereas for Bim, the FRET levels were sustained (Figure 3.4A and SOM Section 5.1). This phenomenon is quantified in Figure 3.4C, which shows the difference between the maximum and endpoint FRET values (due to the effects of noise on estimates of the maximum value, fitted curves were used for these calculations; Methods). The differences in FRET dynamics between DAC-cBid and DAC-Bim suggest that in the later phases of activation, these proteins either adopt different bound states with Bax or have differential affinities for Bax pre- or post-activation. Nevertheless, the fact that both activators have substantial FRET with at least a subset of Bax residues at the experimental endpoint indicates that they remain associated with Bax at a point when most Bax is in pores, consistent with our previous findings for NBD-126C-Bax (Lovell et al., 2008).

3.2.6 THE INITIAL CONFORMATIONAL TRANSITION OF BAX CORRELATES WITH ACTIVATOR:BAX COMPLEX FORMATION, THE LATTER TO BAX OLIGOMERIZATION

To determine the extent to which the initial conformational change in Bax indicated by NBD fluorescence coincided with the formation of the BH₃-only/Bax complex, we fit the NBD and FRET timecourse data




Figure 3.4: FRET measurements of DAC-83C-cBid and DAC-89C-Bim-L with Bax.

(A) Simultaneous measurements of dye release, NBD-Bax fluorescence, and BH3-only/Bax FRET for NBD-54C- and NBD-126C-Bax.

(B) Endpoint FRET between BH3-only activators and labeled Bax mutants.

(C) Difference between maximal and endpoint BH3-only/Bax FRET.

(D) Joint fits of three-conformation model to NBD fluorescence and DAC-cBid:NBD-Bax FRET timecourses.

(E) Joint fits of three-conformation model to NBD fluorescence and DAC-Bax:NBD-Bax FRET timecourses.

simultaneously with the three-conformation model, using two sets of fluorescence scaling parameters: two parameters for the NBD fluorescence of the intermediate and final states, and two parameters for the BH₃-only:Bax FRET. We found that this model fit the data for NBD-54C-Bax and NBD-126C-Bax well using only a single set of rate parameters (Figure 3.4D). The marginal posterior distributions for the FRET of the intermediate and final states indicated that the change in FRET occurred primarily in the $\text{Bax}_{F1} \rightarrow \text{Bax}_{F2}$ transition, declining somewhat in the $\text{Bax}_{F2} \rightarrow \text{Bax}_{F3}$ transition.

To further identify the relationship between the Bax's conformational changes and Bax dimerization, we collected Bax-Bax FRET data for two NBD-Bax mutants representing the the BH₃ and pore-forming regions, with DAC-126C-Bax as the fluorescence donor (Lovell et al., 2008; Kale et al., 2014) (plots of raw data in SOM Section 6.1). The data suggest that unlike FRET with cBid and Bim, Bax-Bax FRET increases fairly late relative to the conformational changes of Bax. Using the same analysis as for the BH₃-only FRET data, we found that a three-conformation model with a single set of rate parameters could simultaneously fit both the NBD and Bax-Bax FRET data (Figure 3.4E). In addition, the parameter estimates indicated that the majority of the increase in Bax-Bax FRET occurred in the latter Bax conformational change ($\text{Bax}_{F2} \rightarrow \text{Bax}_{F3}$).

3.2.7 CANCER-ASSOCIATED BAX MUTANTS HAVE DEFECTS IN SPECIFIC CONFORMATIONAL TRANSITIONS REQUIRED FOR PORE FORMATION

To determine the functional significance of the conformational intermediate we identified, we assessed a series of Bax mutants with diminished pore formation activity to determine whether apoptotic blocks existed primarily at the initial transition, final transition, or both. Because failure to commit apoptosis is considered a hallmark of cancer (Hanahan and Weinberg, 2011), we hypothesized that cancer-associated Bax mutations would be enriched for those lacking apoptotic activity. We therefore generated Bax mutants reported in genome sequencing studies of cancer patients and cell lines. While the most frequent mutation was a frameshift deletion upstream of the BH₃ region, there were also a number of relatively infrequent point mutations reported throughout the BH₃ and pore forming regions (Figure 3.5A).

We cloned and purified 19 of these mutants and analyzed their pore forming activity using an ANTS/DPX




Figure 3.5: Mechanistic characterization of Bax point mutations in cancer.

(A) Genomic alterations in Bax identified by cancer genome sequencing.

(B) ANTS/DPX release activity of Bax point mutants.

(C) Induction of Smac-mCherry release from isolated mitochondria by Bax point mutants.

(D) Endpoint dye release, cBid-Bax FRET, and relative NBD-126C fluorescence of labeled Bax point mutations on a 126C-Bax background.

dye release assay (Figure 3.5A; Methods). As a measure of mutant activity, we determined the EC₅₀ value associated with 50% release at the timecourse endpoint (Methods), and compared it with the EC₅₀ value associated with wild-type Bax. Interestingly, by this measure 10 of the 19 mutants had equal or slightly *increased* activity relative to wild-type Bax, suggesting that their selective effect on tumorigenesis, if any, was independent of Bax's pore forming activity. Of the remaining nine mutants with diminished activity, six of these (G179E, G67R, L113P, S184E, G108E, and G108V) had strongly diminished activity, requiring at least twice as much Bax to achieve wild-type levels of release.

To characterize the dye-release deficient mutants in a more physiologically relevant setting, we assayed their ability to release Smac-mCherry from isolated mitochondria (Methods). The results were similar to the dye release assay, with the mutants showing significantly impaired release activity relative to wild-type Bax (Figure 3.5B). While the degree of inhibition of Smac-mCherry release tended to correspond to the degree of inhibition of dye release, the G179E mutant was a notable exception with relatively minor inhibition of dye release but strong inhibition of Smac-mCherry release. This suggests that either the effect of the G179E mutation is dependent on the proteolipid constituents of the mitochondrial membrane, or that the mutant may be able to form pores small enough to release small molecules but not larger proteins like Smac-mCherry.

To identify the basis for the failure of pore formation in these six mutants, we prepared a corresponding set of mutants labeled at the $\alpha 6$ helix residue C126 and characterized their insertion and pore formation using our combined spectroscopy assay (Figure 3.5C). Two of these mutants, G108E and L113P, showed little activity across all three fluorescence measures, with minimal pore formation, very little FRET with DAC-83C-cBid, and reduced NBD-126C-Bax fluorescence relative to NBD-126C-Bax. The lack of activity in both upstream (FRET with cBid) and downstream (pore formation) measures suggests that these mutants were blocked in the initial conformational transition, $Bax_{F1} \rightarrow Bax_{F2}$. However, three of the mutants showed evidence of being blocked in the intermediate conformation due to a failure of the $Bax_{F2} \rightarrow Bax_{F3}$ transition, with elevated cBid-Bax fluorescence coupled to minimal $\alpha 6$ -helix insertion and pore formation. Finally, the G179E mutant showed an unusual pattern of activity, with pore formation activity nearing that of wild-type but relatively low $\alpha 6$ -helix insertion and elevated Bid-Bax FRET.

3.3 DISCUSSION

A key challenge in understanding the pore formation mechanism of Bax has been the difficulty of obtaining a dynamic picture of the behavior of the full-length protein in membranes at high structural resolution. In this study we combine an extensive dataset of kinetic measurements of fluorescently labeled Bax mutants with an analytical framework that allows us to integrate data from these mutants despite their differing activities. Taken together, our results indicate that Bax undergoes at least two significant structural rearrangements en route to dimerization and pore formation, passing through an intermediate state associated with activator binding and significant conformational changes across the entire protein. The transition from this intermediate to the final state is concomitant with Bax homo-oligomerization.

We identified the number of relevant conformational states by calibrating an ensemble of models to our experimental data. While the two-transition, three-conformation model captured the dynamics of virtually all of the labeled mutants we tested, data from several mutants suggested a high probability of a fourth conformation (third transition). Though Bayes factor calculations indicate that for some mutants the four conformation model is many times more likely than three, the size of the effect is likely to be small—that is, the additional conformation is either at relatively low concentration or involves only a very slight conformational change. One interpretation is that this fourth conformation involves subtle environmental changes that occur during large-scale oligomerization or pore enlargement; another possibility is that the apparent support for a fourth conformation is due to kinetic phenomena that are outside the scope of our class of simple models (saturation of BH₃-only activators, bimolecular reactions, etc.). Proper attribution of this additional timescale evident in the Bax activation process will require additional modeling and experiments.

Despite the high temporal resolution and low noise of many of our experimental timecourses, and the relatively low dimensionality of our kinetic models, we nevertheless found that there was substantial parametric non-identifiability in the fitted parameters, often taking the form of bimodality in the posterior distributions. This highlights the importance of accounting for parametric non-identifiability in modeling studies, either through Bayesian or likelihood-based methods (Kreutz et al., 2013; Eydgahi et al., 2013).

There have been a handful of previous studies aimed at understanding the Bax pore formation mechanism via kinetic modeling (Saito et al., 2000; Kushnareva et al., 2012). In one of the most recent, Kushnareva et al. identified a “lag phase” in pore formation that was apparent in dye release measurements using mitochondrial Outer Membrane Vesicles (OMVs) but not synthetic liposomes (Kushnareva et al., 2012). Our results suggest that there is a lag phase in dye release for synthetic liposomes as well, though it is far less pronounced than what is observed for OMVs, and that it is associated with the time required for Bax to pass through the intermediate conformation en route to forming a fully assembled pore.

It is notable that while engagement of the hydrophobic groove of Bax by a BH₃-only activator activates it (Czabotar et al., 2013; Dai et al., 2011; Moldoveanu et al., 2013), binding of this groove by the BH₃ region of another Bax molecule is necessary for pore formation (Czabotar et al., 2013; Dewson et al., 2008; Dewson et al., 2012). The fact that the addition of excess BH₃-only activator does not appear to inhibit Bax suggests that the hydrophobic groove changes conformation in such a way as to allow Bax to discriminate between activator binding and dimerization. It is possible that this rearrangement of the BH₃ region and hydrophobic groove is the basis of the intermediate conformation we identify in our study.

There has been recent interest in developing inhibitors of Bax and Bak as a means to prevent pathologic cell death in clinical applications, especially as inhibition of executioner caspases alone has proven insufficient to halt death after MOMP (Galluzzi et al., 2014). The dual nature of the hydrophobic groove of Bax as both a site of activation and dimerization suggests that it may be difficult to identify ligands that are primarily inhibitors and not agonists. However, our results suggest that the hydrophobic groove of the conformational intermediate state may be a good target site for an inhibitor, as it could bind specifically to this partially-activated conformation to prevent Bax dimerization and pore formation.

In our study we functionally assessed a series of Bax point mutants derived from cancer genome sequencing studies and found that some mutants appeared to have more pore formation activity than wild-type Bax, and others less. Notably, by far the most frequent genomic alteration of Bax in cancer in the TCGA dataset is a frameshift deletion near the N-terminus, supporting the view that loss of Bax pore formation activity is in most cases tumor-enhancing. While the functional role of the less common point mutations in altering tumor growth and survival is more difficult to interpret, it is important to

point out that Bax conformations that are blocked at the intermediate conformation (e.g., G67R, G108V, and S184E) not only cannot form pores themselves but may also inhibit the activation of wild-type Bax and/or Bak by binding and sequestering BH₃-only activators. However, the magnitude of this effect for the point mutants remains to be characterized.

While this study accounts for the kinetics of the changes in Bax at a number of distinct locations along the protein, it does so primarily for a single set of concentrations of Bax, membranes, and activator. For a clearer picture of the dynamics of pore formation by Bax and the dependence of these dynamics on the balance of the various players, additional kinetic studies will be needed, employing titration of the various components.

3.4 METHODS

3.4.1 PROTEIN PREPARATION AND FLUORESCENT SPECTROSCOPY

Procedures for purifying and labeling Bcl-2 family proteins, preparing lipid vesicles, and making kinetic measurements by fluorescence spectroscopy were described in detail previously ([Kale et al., 2014](#)).

3.4.2 MITOCHONDRIAL SMAC-MCHERRY RELEASE ASSAYS


Mitochondrial permeabilization assays were performed as previously described ([Shamas-Din et al., 2014](#)).

3.4.3 KINETIC MODELING AND SIMULATION

Kinetic models were built programmatically using PySB ([Lopez et al., 2013](#)). The models were formulated as sets of ordinary differential equations and simulated either by 1) numerical integration using the VODE integrator ([Brown et al., 1989](#)) accessed via the Scipy library in Python ([Oliphant, 2007](#)) or 2) solved directly using a closed-form, analytical solution for the system.

The irreversible transition model with n states has $n - 1$ fluorescence transitions and $2n - 2$ free parameters: $n - 1$ transition rate parameters and $n - 1$ fluorescence scaling parameters; the fluorescence associated with the initial aqueous state of Bax, Bax_{F1} , can be calculated directly from the data.

In the simplest case, a labeled residue undergoes a single, irreversible environmental transition between two states with differential NBD fluorescence, denoted F_1 and F_2 :



In such a scheme, in which the transition between the collections of underlying chemical species associated with the fluorescent states F_1 and F_2 occurs by a single transition step, the change in the overall fluorescence is given by the following equation:

$$NBD(t) = C_1[\text{Bax}_{F_1}(t)] + C_2[\text{Bax}_{F_2}(t)] \quad (3.5)$$

where C_1 and C_2 are parameters indicating the NBD fluorescence of Bax_{F_1} and Bax_{F_2} , respectively. This equation has the solution:

$$NBD(t) = [\text{Bax}_o] (C_2 + (C_1 - C_2)e^{-kt}) \quad (3.6)$$

Since our data is normalized in terms of the fluorescence relative to the initial state Bax_{F_1} , with the assumption that all Bax is in conformation Bax_{F_1} at $t = 0$ (that is, $NBD(0) = \text{Bax}_o C_1$), we can also express this as a two-parameter, single exponential equation:

$$NBD_{norm}(t) = NBD(t) / (\text{Bax}_o C_1) \quad (3.7)$$

$$NBD_{norm}(t) = 1 + F_{max}(1 - e^{-kt}) \quad (3.8)$$

where $F_{max} = \frac{C_2}{C_1} - 1$.

3.4.4 PARAMETER ESTIMATION AND MODEL DISCRIMINATION

Parameter estimation and model discrimination were performed using affine-invariant Markov chain Monte Carlo (MCMC) sampling implemented by the Python software package *emcee* (Foreman-Mackey et al., 2013) and based on the algorithm described by Goodman and Weare, 2010. Parallel tempering was used to aid convergence and to calculate marginal likelihood values by thermodynamic integration

(Geyer, 1991; Lartillot and Philippe, 2006). For each MCMC run, a ladder of 50 different temperatures were used, with 400 walkers at each temperature. Values for the reciprocal temperature, $\beta = 1/T$, were spaced logarithmically from 1 to a maximum temperature of $\beta = 10^{-6}$. Starting positions for the walkers were chosen randomly from the prior distributions for each parameter.

Convergence of the MCMC chains was assessed by several heuristics: 1) the log marginal likelihood ($\log(\text{ML})$) values were calculated by thermodynamic integration every 50 steps and assessed for asymptotic convergence by comparing $\log(\text{ML})$ from the last 50 steps to the $\log(\text{ML})$ value calculated from the previous 50 steps. If the difference in $\log(\text{ML})$ was greater than an absolute threshold of 3, or greater than a relative threshold of $0.1 \times \text{err}$, where err was the error associated with the thermodynamic integration procedure itself, then the chain was considered to be non-convergent. 2) If the chain passed the test of convergence described in (1), the chains at each of the 50 temperatures were assessed for any trend towards increases in posterior probability over the 50-step convergence interval by performing linear regression on the posterior probability values associated with the sampled positions. If the trend for any of the 50 chains was positive with p-value less than 0.001, the chain ensemble was considered to be non-convergent. If a chain passed the programmatic heuristics in 1) and 2), the “burn-in” period was terminated and samples were recorded for an additional 100 steps, yielding $100 \text{ steps} \times 400 \text{ walkers} = 40,000$ samples for each parameter at each temperature. Chains were assessed for mixing by inspection of the posteriors associated with the positions of each walker across a subset of temperatures, which are plotted in **SOM Section 3.4**. The frequency of accepted temperature swaps was also inspected to ensure proper mixing between chains at different temperatures. Sampling runs were performed primarily on a computing cluster assembled from Amazon Elastic Compute Cloud (EC2) instances using the StarCluster software package (Fusaro et al., 2011). All code necessary to reproduce the results of data analysis and model calibration is freely available on GitHub, at <https://github.com/johnbachman/tBidBaxLipo>.

3.4.5 CALCULATION OF MAXIMUM BH₃-ONLY/BAX FRET VALUES

FRET timecourses for Bid and Bim with Bax were fitted with the three-conformation, irreversible transition model to get a continuous curve approximating the fluorescence data. The maximum FRET value

was then calculated from the fitted curve.

3.5 SUPPLEMENTARY INFORMATION

Many chemical processes can be described as a series of linked equilibria:



A common practice in biochemistry for studying such processes is to prepare an *in vitro* system with F in a well-defined initial state and monitor its evolution towards equilibrium to determine the number of relevant states and the timescales associated with their associated transitions. These systems are often monitored experimentally using fluorescence spectroscopy with labeled components having differential fluorescence intensities for different states. Examples include measurements of molecular interactions by FRET, protein-membrane interactions using the environmentally sensitive label NBD, protein-nucleotide interactions using labeled nucleotides, etc. Given that such experimental approaches only allow the monitoring of the total fluorescence, not the fluorescences or concentrations of the intermediates, mathematical models are often necessary to make inferences about the system.

In this analysis we consider three classes of mathematical models, increasingly mechanistic in their approach, that can be used to fit data from these types of experiments:

1. a sum of $n - 1$ exponentials,
2. a system of ordinary differential equations for a set of *irreversible* transitions from F_1 to F_n , or
3. a system of ODEs for *reversible* transitions between F_1 and F_n .

We show that, in terms of their abilities to fit the observed fluorescence data generated by the physical process, and to identify the *number* of conformational states, these three classes of models are indistinguishable: their solutions all take the form of a sum of $n - 1$ exponentials.

3.5.1 EXPERIMENTAL SETUP AND DEFINITIONS

We consider systems in which the total amount of the labeled entity F is constant during the experiment, indicated by the conservation equation

$$F_{tot} = \sum_{i=1}^n F_i(t) \quad (3.10)$$

The distinct states F_i have fluorescence intensities defined by parameters C_i , giving the following expression for the experimentally observed (total) fluorescence, $E(t)$:

$$E(t) = \sum_{i=1}^n C_i[F_i(t)] \quad (3.11)$$

with $[F_i(t)]$ denoting the concentration of the species F_i over time.

SUM OF EXPONENTIALS

The most straightforward and commonly-used approach to analyzing equilibrium systems is to fit a sum of exponentials such as the following:

$$E(t) = \mathcal{A}_1 e^{-k_1 t} + \mathcal{A}_2 e^{-k_2 t} + \dots + \mathcal{A}_n e^{-k_n t} + C \quad (3.12)$$

where the number of exponentials required to fit the data is indicative of the number of relevant timescales. A model consisting of n exponentials has $2n + 1$ free parameters: n scaling coefficients \mathcal{A}_i , n time constants k_i and the constant factor C (which may be required to account for background fluorescence).

IRREVERSIBLE TRANSITION MODEL

The system of first-order, irreversible transitions between n states is defined $n - 1$ rate parameters:



It is defined by the system of ODEs


$$\frac{dF_1}{dt} = -k_1 F_1 \quad (3.14)$$

$$\frac{dF_i}{dt} = k_{i-1} F_{i-1} - k_i F_i \quad (3.15)$$

$$\frac{dF_n}{dt} = k_{n-1} F_{n-1} \quad (3.16)$$

REVERSIBLE TRANSITION MODEL

The system of linked equilibria with n states has $n - 1$ transitions defined by corresponding forward and reverse parameters:



This system is defined by the system of ODEs

$$\frac{dF_1}{dt} = -k_1^F F_1 + k_1^R F_2 \quad (3.18)$$

$$\frac{dF_i}{dt} = k_{i-1}^F F_{i-1} - k_{i-1}^R F_i - k_i^F F_i + k_i^R F_{i+1} \quad (3.19)$$

$$\frac{dF_n}{dt} = k_{n-1}^F F_{n-1} - k_{n-1}^R F_n \quad (3.20)$$

3.5.2 GENERAL SOLUTION FOR THE IRREVERSIBLE MODEL

The system of ODEs for the irreversible model can be described in matrix form as

$$\begin{bmatrix} F'_1 \\ F'_2 \\ \vdots \\ F'_{n-1} \\ F'_n \end{bmatrix} = \begin{bmatrix} -k_1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ k_1 & -k_2 & 0 & \cdots & 0 & 0 & 0 \\ \vdots & & & \ddots & & & \\ 0 & 0 & 0 & \cdots & k_{n-2} & -k_{n-1} & 0 \\ 0 & 0 & 0 & \cdots & 0 & k_{n-1} & 0 \end{bmatrix} \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_{n-1} \\ F_n \end{bmatrix} \quad (3.21)$$

Because the final state F_n does not appear on the right hand side of any of the equations, it can be removed to obtain an equivalent system for the first $n - 1$ states:

$$\begin{bmatrix} F'_1 \\ F'_2 \\ \vdots \\ F'_{n-1} \end{bmatrix} = \begin{bmatrix} -k_1 & 0 & 0 & \cdots & 0 & 0 \\ k_1 & -k_2 & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & & \\ 0 & 0 & 0 & \cdots & k_{n-2} & -k_{n-1} \end{bmatrix} \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_{n-1} \end{bmatrix} \quad (3.22)$$

After obtaining the solution to the reduced system in (3.22), the concentration of the final state F_n can be calculated from the conservation equation

$$F_n = F_{tot} - \sum_{i=1}^{n-1} F_i \quad (3.23)$$

The general solution of (3.22) takes the form

$$\vec{F}(t) = \sum_{j=1}^{n-1} a_j e^{\lambda_j t} \vec{\eta}^{(j)} \quad (3.24)$$

where the a_j represent arbitrary constants and λ_j and $\vec{\eta}^{(j)}$ denote the j th eigenvalue and eigenvector of \mathcal{A} , respectively.

The individual entries in $\vec{F}(t)$ are therefore given by

$$F_i(t) = \sum_{j=1}^{n-1} a_j e^{\lambda_j t} \vec{\eta}_i^{(j)}, \quad 1 \leq j \leq n-1 \quad (3.25)$$

where $\vec{\eta}_i^{(j)}$ denotes the i th element of the j th eigenvector of the matrix A .

By applying the conservation equation in (3.23) to the equation for the fluorescence (3.11), we get an expression for the observed fluorescence in terms of the first $n-1$ states:

$$E(t) = \sum_{i=1}^{n-1} C_i F_i(t) + C_n \left(F_{tot} - \sum_{i=1}^{n-1} F_i \right) \quad (3.26)$$

$$= \sum_{i=1}^{n-1} (C_i - C_n) F_i(t) + C_n F_{tot} \quad (3.27)$$

Substituting in for $\vec{F}(t)$ we obtain the general equation for the fluorescence for the irreversible model with n states:

$$E(t) = \sum_{i=1}^{n-1} (C_i - C_n) F_i(t) + C_n F_{tot} \quad (3.28)$$

$$= \sum_{i=1}^{n-1} \left((C_i - C_n) \sum_{j=1}^{n-1} a_j e^{\lambda_j t} \vec{\eta}_i^{(j)} \right) + C_n F_{tot} \quad (3.29)$$

$$= \sum_{j=1}^{n-1} \left(e^{\lambda_j t} \sum_{i=1}^{n-1} (C_i - C_n) a_j \vec{\eta}_i^{(j)} \right) + C_n F_{tot} \quad (3.30)$$

If we define

$$\alpha_j = \sum_{i=1}^{n-1} (C_i - C_n) a_j \vec{\eta}_i^{(j)} \quad (3.31)$$

we get our final expression

$$E(t) = \sum_{j=1}^{n-1} \alpha_j e^{\lambda_j t} + C_n F_{tot} \quad (3.32)$$

Therefore, the dynamics of a linked equilibrium system is characterized by a sum of exponentials defining the timescales. In general, a system with n states and $n - 1$ transitions will consist of a sum of $n - 1$ exponentials. Situations that will reduce the apparent number of timescales include:

- If any of the α_j are 0
- If any of the eigenvalues of the matrix are 0
- If any of the eigenvalues are equal (occurs more than once)

3.5.3 THE TIMESCALES OF THE IRREVERSIBLE MODEL ARE DEFINED BY THE RATE CONSTANTS k_i

As shown by the general form of the irreversible model (3.32), the fluorescence is given by the sum of exponentials with the time constants for each exponential determined by the eigenvalues of the matrix \mathcal{A} .

The eigenvalues of \mathcal{A} are given by the solutions to the characteristic polynomial

$$p_{\mathcal{A}}(\lambda) = \det(\lambda I - \mathcal{A}) \quad (3.33)$$

Because the characteristic matrix $\lambda I - \mathcal{A}$ is lower triangular, its determinant is the product of the diagonal entries:

$$p_{\mathcal{A}}(\lambda) = (\lambda + k_1)(\lambda + k_2) \cdots (\lambda + k_{n-1}) \quad (3.34)$$

with roots $k_i, 1 \leq i \leq n - 1$.

3.5.4 GENERAL SOLUTION FOR THE REVERSIBLE MODEL

The system of ODEs for the reversible model can be described in matrix form as

$$\begin{bmatrix} F_1' \\ F_2' \\ \vdots \\ F_{n-1}' \\ F_n' \end{bmatrix} = \begin{bmatrix} -k_1^F & k_1^R & 0 & 0 & \cdots & 0 & 0 & 0 \\ k_1^F & -(k_1^R + k_2^F) & k_3^R & \cdots & 0 & 0 & 0 \\ \vdots & & & \ddots & & & \\ 0 & 0 & 0 & \cdots & k_{n-2}^F & -(k_{n-2}^R + k_{n-1}^F) & k_{n-1}^R \\ 0 & 0 & 0 & \cdots & 0 & k_{n-1}^F & -k_{n-1}^R \end{bmatrix} \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_{n-1} \\ F_n \end{bmatrix} \quad (3.35)$$

By applying the conservation equation (3.23), this can be reduced to a system of $n - 1$ equations:

$$\begin{bmatrix} F_1' \\ F_2' \\ \vdots \\ F_{n-1}' \end{bmatrix} = \begin{bmatrix} -k_1^F & k_1^R & 0 & \cdots & 0 & 0 \\ k_1^F & -(k_1^R + k_2^F) & k_3^R & \cdots & 0 & 0 \\ \vdots & & & \ddots & & \\ -k_{n-1}^R & -k_{n-1}^R & -k_{n-1}^R & \cdots & k_{n-2}^F - k_{n-1}^R & -(k_{n-2}^R + k_{n-1}^F + k_{n-1}^R) \end{bmatrix} \begin{bmatrix} F_1 \\ F_2 \\ \vdots \\ F_{n-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ k_{n-1}^R F_{tot} \end{bmatrix} \quad (3.36)$$

As for the irreversible system, the concentration of the final state F_n can be calculated from the conservation equation after solving the reduced system (3.36).

As with the irreversible model, solution of this system yields a sum of $n - 1$ exponentials plus a constant term, indicating that the three types of models will produce equivalent fits to the data despite their different parameterizations.

References

- Albeck, J. G., J. M. Burke, S. L. Spencer, D. A. Lauffenburger, and P. K. Sorger (2008). “Modeling a snap-action, variable-delay switch controlling extrinsic cell death”. *PLoS Biology* 6.12 (Dec. 2008), pp. 2831–2852.
- Annis, M. G., E. L. Soucie, P. J. Dlugosz, J. A. Cruz-Aguado, L. Z. Penn, B. Leber, and D. W. Andrews (2005). “Bax forms multispinning monomers that oligomerize to permeabilize membranes during apoptosis”. *The EMBO Journal* 24.12 (June 2005), pp. 2096–2103.
- Bleicken, S., G. Jeschke, C. Stegmüller, R. Salvador-Gallego, A. J. García-Sáez, and E. Bordinon (2014). “Structural model of active Bax at the membrane.” *Molecular Cell* 56.4 (Nov. 2014), pp. 496–505.
- Brown, P., G. Byrne, and A. C. Hindmarsh (1989). “VODE, A variable coefficient ODE solver”. *SIAM J Sci Stat Comput* 10, pp. 1038–1051.
- Certo, M., V. Del Gaizo Moore, M. Nishino, G. Wei, S. Korsmeyer, S. A. Armstrong, and A. Letai (2006). “Mitochondria primed by death signals determine cellular addiction to antiapoptotic BCL-2 family members”. *Cancer Cell* 9.5 (May 2006), pp. 351–365.
- Chipuk, J. E., T. Moldoveanu, F. Llambi, M. J. Parsons, and D. R. Green (2010). “The BCL-2 family reunion.” *Molecular Cell* 37.3 (Feb. 2010), pp. 299–310.
- Chonghaile, T. N. et al. (2011). “Pretreatment Mitochondrial Priming Correlates with Clinical Response to Cytotoxic Chemotherapy”. *Science (New York, NY)* 334.6059 (Nov. 2011), pp. 1129–1133.
- Czabotar, P. E. et al. (2013). “Bax crystal structures reveal how BH₃ domains activate Bax and nucleate its oligomerization to induce apoptosis.” *Cell* 152.3 (Jan. 2013), pp. 519–531.
- Dai, H., A. Smith, X. W. Meng, P. A. Schneider, Y.-P. Pang, and S. H. Kaufmann (2011). “Transient binding of an activator BH₃ domain to the Bak BH₃-binding groove initiates Bak oligomerization.” *The Journal of Cell Biology* 194.1 (July 2011), pp. 39–48.
- Dewson, G., S. Ma, P. Frederick, C. Hockings, I. Tan, T. Kratina, and R. M. Kluck (2012). “Bax dimerizes via a symmetric BH₃:groove interface during apoptosis.” *Cell Death and Differentiation* 19.4 (Apr. 2012), pp. 661–670.
- Dewson, G., T. Kratina, H. W. Sim, H. Puthalakath, J. M. Adams, P. M. Colman, and R. M. Kluck (2008). “To trigger apoptosis, Bak exposes its BH₃ domain and homodimerizes via BH₃:groove interactions.” *Molecular Cell* 30.3 (May 2008), pp. 369–380.
- Eydgahi, H., W. W. Chen, J. L. Muhlich, D. Vitkup, J. N. Tsitsiklis, and P. K. Sorger (2013). “Properties of cell death models calibrated and compared using Bayesian approaches”. *Molecular Systems Biology* 9.1 (Jan. 2013), pp. 644–644.

- Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman (2013). “emcee: The MCMC Hammer”. *Publications of the Astronomical Society of the Pacific* 125.925 (Mar. 2013), pp. 306–312.
- Fusaro, V. A., P. Patil, E. Gafni, D. P. Wall, and P. J. Tonellato (2011). “Biomedical cloud computing with Amazon Web Services.” *PLoS Computational Biology* 7.8 (Aug. 2011), e1002147.
- Galluzzi, L. et al. (2014). “Essential versus accessory aspects of cell death: recommendations of the NCCD 2015.” *Cell Death and Differentiation* (Sept. 2014).
- Gavathiotis, E., D. E. Reyna, J. A. Bellairs, E. S. Leshchiner, and L. D. Walensky (2012). “Direct and selective small-molecule activation of proapoptotic BAX.” *Nature Chemical Biology* 8.7 (July 2012), pp. 639–645.
- Gavathiotis, E., D. E. Reyna, M. L. Davis, G. H. Bird, and L. D. Walensky (2010). “BH₃-Triggered Structural Reorganization Drives the Activation of Proapoptotic BAX”. *Molecular Cell* 40.3 (Nov. 2010), pp. 481–492.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (2014). *Bayesian data analysis*. Chapman and Hall/CRC.
- George, N. M., J. J. D. Evans, and X. Luo (2007). “A three-helix homo-oligomerization domain containing BH₃ and BH₁ is responsible for the apoptotic activity of Bax”. *Genes & Development* 21.15 (Aug. 2007), pp. 1937–1948.
- Geyer, C. J. (1991). “Markov chain Monte Carlo maximum likelihood”. *Computing Science and Statistics: Proc. 23rd Symposium on the Interface, Interface Foundation, Fairfax Station, VA*, pp. 156–163.
- Girolami, M., B. Calderhead, and V. Vyshemirsky (2010). “System Identification and Model Ranking: The Bayesian Perspective Learning and Inference”. In: *Learning and Inference in Computational Systems Biology*. MIT Press.
- Goldstein, J. C., N. J. Waterhouse, P. Juin, G. I. Evan, and D. R. Green (2000). “The coordinate release of cytochrome c during apoptosis is rapid, complete and kinetically invariant”. *Nature Cell Biology* 2.3 (Mar. 2000), pp. 156–162.
- Goodman, J. and J. Weare (2010). “Ensemble samplers with affine invariance”. *Communications in Applied Mathematics and Computational Science* 5.1 (Jan. 2010), pp. 65–80.
- Gunawardena, J. (2012). “Some lessons about models from Michaelis and Menten.” *Molecular Biology of the Cell* 23.4 (Feb. 2012), pp. 517–519.
- Hanahan, D. and R. A. Weinberg (2011). “Hallmarks of cancer: the next generation.” *Cell* 144.5 (Mar. 2011), pp. 646–674.
- Kale, J., X. Chi, B. Leber, and D. Andrews (2014). “Examining the molecular mechanism of bcl-2 family proteins at membranes by fluorescence spectroscopy.” *Methods in enzymology* 544, pp. 1–23.

- Kim, H., H.-C. Tu, D. Ren, O. Takeuchi, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2009). "Stepwise activation of BAX and BAK by tBID, BIM, and PUMA initiates mitochondrial apoptosis". *Molecular Cell* 36.3 (Nov. 2009), pp. 487–499.
- Kreutz, C., A. Raue, D. Kaschek, and J. Timmer (2013). "Profile likelihood in systems biology." *The FEBS journal* 280.11 (June 2013), pp. 2564–2571.
- Kushnareva, Y., A. Y. Andreyev, T. Kuwana, and D. D. Newmeyer (2012). "Bax activation initiates the assembly of a multimeric catalyst that facilitates Bax pore formation in mitochondrial outer membranes." *PLoS Biology* 10.9, e1001394.
- Lartillot, N. and H. Philippe (2006). "Computing Bayes factors using thermodynamic integration." *Systematic biology* 55.2 (Apr. 2006), pp. 195–207.
- Letai, A., M. C. Bassik, L. D. Walensky, M. D. Sorcinelli, S. Weiler, and S. J. Korsmeyer (2002). "Distinct BH3 domains either sensitize or activate mitochondrial apoptosis, serving as prototype cancer therapeutics". *Cancer Cell* 2.3 (Sept. 2002), pp. 183–192.
- Levenson, J. D. et al. (2015). "Potent and selective small-molecule MCL-1 inhibitors demonstrate on-target cancer cell killing activity as single agents and in combination with ABT-263 (navitoclax)." *Cell death & disease* 6, e1590.
- Lopez, C. F., J. L. Muhlich, J. A. Bachman, and P. K. Sorger (2013). "Programming biological models in Python using PySB." *Molecular Systems Biology* 9, p. 646.
- Lovell, J. F., L. P. Billen, S. Bindner, A. Shamas-Din, C. Fradin, B. Leber, and D. W. Andrews (2008). "Membrane binding by tBid initiates an ordered series of events culminating in membrane permeabilization by Bax". *Cell* 135.6 (Dec. 2008), pp. 1074–1084.
- Moldoveanu, T., C. R. Grace, F. Llambi, A. Nourse, P. Fitzgerald, K. Gehring, R. W. Kriwacki, and D. R. Green (2013). "BID-induced structural changes in BAK promote apoptosis." *Nature Structural & Molecular Biology* 20.5 (May 2013), pp. 589–597.
- Oliphant, T. E. (2007). "Python for Scientific Computing". *Computing in Science & Engineering* 9.3, pp. 10–20.
- Oltersdorf, T. et al. (2005). "An inhibitor of Bcl-2 family proteins induces regression of solid tumours". *Nature* 435.7042 (June 2005), pp. 677–681.
- Rehm, M., H. Dussmann, R. U. Janicke, J. M. Tavaré, D. Kogel, and J. H. M. Prehn (2002). "Single-cell fluorescence resonance energy transfer analysis demonstrates that caspase activation during apoptosis is a rapid process. Role of caspase-3". *The Journal of biological chemistry* 277.27 (July 2002), pp. 24506–24514.
- Saito, M., S. J. Korsmeyer, and P. H. Schlesinger (2000). "BAX-dependent transport of cytochrome c reconstituted in pure liposomes." *Nature Cell Biology* 2.8 (Aug. 2000), pp. 553–555.

- Shamas-Din, A., D. Satsoura, O. Khan, W. Zhu, B. Leber, C. Fradin, and D. W. Andrews (2014). "Multiple partners can kiss-and-run: Bax transfers between multiple membranes and permeabilizes those primed by tBid." *Cell death & disease* 5, e1277.
- Souers, A. J. et al. (2013). "ABT-199, a potent and selective BCL-2 inhibitor, achieves antitumor activity while sparing platelets." *Nature medicine* 19.2 (Feb. 2013), pp. 202–208.
- Suzuki, M., R. J. Youle, and N. Tjandra (2000). "Structure of Bax: Coregulation of Dimer Formation and Intracellular Localization". *Cell*.
- Tait, S. W. G. and D. R. Green (2010). "Mitochondria and cell death: outer membrane permeabilization and beyond". *Nature Reviews Molecular Cell Biology* 11.9 (Sept. 2010), pp. 621–632.
- Vo, T.-T., J. Ryan, R. Carrasco, D. Neuberg, D. J. Rossi, R. M. Stone, D. J. Deangelo, M. G. Frattini, and A. Letai (2012). "Relative mitochondrial priming of myeloblasts and normal HSCs determines chemotherapeutic success in AML." *Cell* 151.2 (Oct. 2012), pp. 344–355.
- Westphal, D. et al. (2014). "Apoptotic pore formation is associated with in-plane insertion of Bak or Bax central helices into the mitochondrial outer membrane." *Proceedings of the National Academy of Sciences of the United States of America* (Sept. 2014).
- Yethon, J. A., R. F. Epand, B. Leber, R. M. Epand, and D. W. Andrews (2003). "Interaction with a membrane surface triggers a reversible conformational change in Bax normally associated with induction of apoptosis." *The Journal of biological chemistry* 278.49 (Dec. 2003), pp. 48935–48941.

[The blind men said,] “we cannot agree to what the elephant is like.” The wise man then calmly said, “Each one of you is correct; and each one of you is wrong. Because each one of you had only touched a part of the elephant’s body. Thus you only have a partial view of the animal. If you put your partial views together, you will get an idea of what an elephant looks like.”

Jain parable

4

Identifying the insertion kinetics, stability and stoichiometry of the Bax pore from fluorescence measurements and mathematical modeling

4.1	Introduction	122
4.2	Results	127
4.3	Discussion	146
4.4	Methods	149

AUTHORSHIP CONTRIBUTION

For this study I performed all experiments and data analysis with the exception of the Alexa 568-cBid/membrane binding study, which was performed by J. Kale. I prepared the figures and wrote the manuscript.

ABSTRACT

In this study we develop an integrated understanding of the dynamics and mechanism of apoptotic pore formation by coupling an *in vitro* biochemical model system with quantitative measurement and systematic mathematical modeling. We measure the Bax recruitment and pore formation kinetics of synthetic lipid vesicles incubated with purified cBid and Bax and use these measurements to distinguish between competing mechanistic hypotheses for Bcl-2 protein function. Mechanistic models are enumerated systematically, calibrated to data, and discriminated by Bayesian statistical methods. The results from this approach highlight several previously unappreciated features of the Bax pore formation mechanism: 1) despite being considered an interaction with fast turnover (“hit and run”), the activation of Bax by cBid is saturable, limiting rates of Bax recruitment; 2) while Bax recruitment to synthetic lipid membranes is not limited at physiologic concentrations, cBid recruitment to membranes is saturable, likely due to unfavorable electrostatic interactions between Bid molecules at the vesicle surface; 3) the Bax pore is stable, with a minimum size of four subunits for the release of small molecules; and 4) Bax pores cooperatively recruit additional Bax subunits, but in a fashion that actually diminishes, rather than increases, the overall number of permeabilized vesicles. More generally, our approach highlights that data drawn from simplified systems carry a great deal of mechanistic information that can be extracted by rigorous model-based inference grounded in prior knowledge of the relevant biochemistry.

4.1 INTRODUCTION

A key regulatory step in apoptosis is the formation of pores in the mitochondrial outer membrane, a process known as MOMP (Tait and Green, 2010). In MOMP, pores in the mitochondrial outer membrane trigger the release of several pro-apoptotic proteins, including cytochrome *c* and Smac, from the mitochondrial intermembrane space into the cytosol, where they trigger the activation of executioner caspases

and subsequently the orderly breakdown of the cell. MOMP itself is regulated by a related set of proteins known as the Bcl-2 family, which includes both pro-survival and pro-apoptotic proteins. Their role as the gatekeepers of programmed cell death has made the Bcl-2 family members key therapeutic targets: most cytotoxic anti-cancer drugs trigger cell death by stimulating the activation of various Bcl-2 family members, and the abundance of pro- and anti-apoptotic Bcl-2 family members at the mitochondrial membrane prior to treatment strongly influences the likelihood of therapeutic success (Chonghaile et al., 2011; Vo et al., 2012). Small molecules have also recently been developed to specifically target Bcl-2 protein family members, by either inhibiting the anti-apoptotic factors Bcl-2, Bcl-XL, and Mcl-1 (Oltersdorf et al., 2005; Levenson et al., 2015; Souers et al., 2013) or activating the pore forming protein Bax (Gavathiotis et al., 2012).

Despite extensive progress in determining the structures of Bcl-2 proteins and their network of interactions, there remain a number of unresolved questions about the pore formation mechanism. An intensively studied but still open question is the structure and stoichiometry of the Bax/Bak pore, which has in the past been termed the “holy grail” of apoptosis research (Volkman et al., 2013). Key questions in this debate are whether the pore edge is lined primarily by Bax proteins (proteinaceous pore) or by lipid headgroups (lipidic pore), and whether its stoichiometry is fixed or variable. Estimates of the minimal stoichiometry of the pore have ranged from two to four (Saito et al., 2000), to nine (Martinez-Caballero et al., 2009). Moreover, imaging and flow cytometry studies have suggested that the Bax pore grows (Blicken et al., 2013; Gillies et al., 2015), implying a variable stoichiometry.

It is unclear how mechanisms that have been shown to regulate Bax activity, including Bax auto-activation and retrotranslocation, relate to pore stoichiometry and stability. The first of these, Bax auto-activation, is a process by which the BH₃ region of activated Bax acts as an activator of other Bax molecules, propagating the apoptotic signal (Tan et al., 2006; Gavathiotis et al., 2010). Mathematical models have suggested that Bax auto-activation plays a key role in establishing an all-or-none mechanism of action in cell death (Cui et al., 2008). However, whether autoactivation is responsible for activating pores at other vesicle sites, creating additional pores on the same vesicle, or enlarging existing pores, is not clear. While Bax has been shown to form large clusters at pore sites by fluorescent imaging and electron microscopy

(Dussman et al., 2009; Nechushtan et al., 2001; Albeck et al., 2008b), it is unknown whether this process is mediated by auto-activation and whether these clusters represent enlarged pores of variable stoichiometry or nonspecific aggregation at the vesicle surface. Since studies demonstrating Bax autoactivation have relied on Bax BH₃ peptides (Tan et al., 2006) or heat (Gavathiotis et al., 2010) to activate Bax, it is also unclear how much this mechanism contributes to overall Bax activation in the presence of an activator such as Bid or Bim.

The role of Bax's translocation equilibrium between cytosol and membranes in regulating the cell's propensity for apoptosis and the stability of the pore also remains poorly understood. Though it has been shown that binding between Bcl-XL and Bax tends to result in an equilibrium with more Bax in the cytosol, a process termed "retrotranslocation" (Billen et al., 2008; Edlich et al., 2011; Todt et al., 2013), what is less clear is at what stage of the activation sequence this occurs. While Bax binds membranes transiently prior to activation (Yethon et al., 2003), it is not known whether Bax readily unbinds membranes after it has undergone activation. This is an important question for MOMP regulation as it determines whether the Bax pore, once formed, is stable: do pores grow inevitably (e.g., by auto-activation) or can they form transiently and then be dismantled upon binding of antiapoptotics such as Bcl-XL? FRET studies have suggested that Bax retrotranslocation from membranes is much slower than the corresponding rate for tBid, suggesting that Bax retrotranslocation plays a relatively minor role after activation and pore formation, at least in the absence of Bcl-XL (Shamas-Din et al., 2014). On the other hand, flow cytometry studies with OMVs have suggested that the Bax pore grows but can also collapse (Gillies et al., 2015).

In addition to the specific mechanistic issues described above, there is a well-recognized need for a more thorough characterization of how the abundance and composition of lipid membranes governs the interactions among the Bcl-2 protein family members. The fact that many Bcl-2 interactions take place only in the membrane means that any assessment of the interactions among these proteins must also take into account their individual interactions with membranes. For example, the interaction of tBid with synthetic membranes was recently shown to be highly dynamic, with a relatively low affinity that is also dependent on the presence of cardiolipin and negatively charged lipid headgroups (Shamas-Din et al., 2015; Shamas-Din et al., 2014; Shamas-Din et al., 2013). It was reported that Bax binding to lipid mem-

branes saturates at fairly low Bax/liposome ratios of 10 (Satsoura et al., 2012), which is surprising given several observations of large pores and clusters of Bax in both *in vitro* and *in vivo* studies (Schafer et al., 2009; Gillies et al., 2015; Nechushtan et al., 2001). Understanding the role of membranes in determining kinetic and stoichiometric constraints on Bax recruitment is essential not only to reproducibly interpret experiments in which lipid concentration is an uncontrolled variable but also to understand the relationship between variability in mitochondrial abundance, morphology, and composition and the propensity of cells to activate Bcl-2 proteins and undergo apoptosis.

Additional mechanistic questions arise from attempts to integrate individually convincing pieces of data from disparate studies. This is illustrated by the ambiguities surrounding the interaction between activator BH₃-only proteins and the pore forming proteins Bax and Bak. Though previously a matter of debate, it is now widely accepted that activator BH₃-only proteins directly bind Bax and Bak to trigger a conformational change that makes them competent for pore formation. NMR and biochemical studies of Bax have identified an initial conformational change involving the ejection of the α_9 helix from the BH_{1:3} groove and the exposure of the 6A7 epitope, mediated by binding of the BH₃-only activator to a “rear pocket” on Bax (Kim et al., 2009; Gavathiotis et al., 2010). Other studies have suggested that the activator binds, possibly in a subsequent step, to the BH_{1:3} groove of Bax itself in a manner analogous to the binding and sequestration of BH₃-only proteins by the anti-apoptotic Bcl-2 family members (Czabotar et al., 2013; Dai et al., 2011).

The interaction between Bax and its BH₃ activators such as Bid and Bim has been found to be transient and is often described as “hit and run.” (K. Wang et al., 1996; Gavathiotis et al., 2010). However, FRET studies with full-length Bid and Bax have shown that Bax remains bound to its activator in the membrane, even after it has inserted into membranes (Lovell et al., 2008); in addition, immunoprecipitation of activated Bax by the 6A7 antibody pulls down tBid (Kim et al., 2009). In light of the known interaction surfaces of Bax and its activators, the role of this post-activation complex presents a conundrum. If the complex involves the BH₃ motif of the activator, the accumulation of active Bax should progressively sequester the activator, slowing down further activation. On the other hand, if the complex with activator involves the BH₃ region or BH_{1:3} groove of Bax, this should inhibit Bax dimerization and pore

formation, which has been shown to involve precisely these surfaces (George et al., 2007; Czabotar et al., 2013; Dewson et al., 2012). This suggests that either the functional effects of this Bax:activator complex have not yet been characterized, or that the post-activation complex between these proteins involves novel binding sites of one or both proteins.

A classic approach to determining reaction mechanisms in biochemical systems is by kinetic analysis of key steps under titration of system components. Kinetic analysis can determine the rate limiting steps under various concentration conditions, the molecularity (e.g., uni- versus bi-) of various reaction steps, and whether binding interactions are competitive or non-competitive. Kinetic analysis has been used extensively to characterize the mechanisms of pore forming peptides and toxins (Schwarz et al., 1992; Gregory et al., 2008). Despite the availability of a well-validated *in vitro* kinetic assay for Bcl-2 protein activity, there have been few detailed modeling studies (Saito et al., 2000; Kushnareva et al., 2012). In addition, despite the demonstrated robustness of the experimental system, these studies have given rise to apparently conflicting results. For example, the relatively straightforward measure of the scaling of the rate of permeabilization with Bax concentration has variously been reported as cooperative (Saito et al., 2000), saturating (Satsoura et al., 2012), and linear (Kushnareva et al., 2012). This may in part be due to difficulties of extending the analytical methods developed for permeabilization by peptides to the Bcl-2 family, which involve three components rather than two (vesicles, pore forming protein, and activator). The addition of a third component not only increases experimental complexity, but also combinatorially increases the number of models to consider and makes parameter estimation and model comparison by classical optimization methods far more challenging.

In this study we measure Bax recruitment and pore formation across a range of experimental conditions and use mathematical models to explain how the results in the full, three-component system depend on the precise location in concentration space. In contrast to models that explain only single slices of this data, such as the response to Bax titration at a single Bid and membrane concentration, the resulting models have explanatory power across the whole concentration space. Moreover, by using this process we were able to resolve questions about the relative roles of Bid, Bax and membranes in determining Bax recruitment and insertion. The data and models show that excess Bax can indeed kinetically inhibit Bid's

activity, but primarily in its inactivated state. We show that membrane insertion of the pore forming helix of Bax occurs concurrently with Bax dimerization and is reversible; that the Bax pore in large unilamellar vesicles is stable, has a minimal size of four subunits, and grows through a process that results in a highly non-independent distribution of Bax across vesicles. Extending our conclusions to MOMP in the biological setting, we propose that Bax auto-activation does not, as often suggested, “propagate” the apoptotic signal across the cell, but rather causes pores to become enlarged locally, which can actually diminish the global response. In this interpretation, the “snap-action” mechanism of apoptosis in cells can be explained by the fact that the Bax-mitochondria system is a high-membrane concentration, low compartment number system, leading to rapid recruitment of Bax and enlargement of pores.

4.2 RESULTS

To identify the mechanism of Bax pore formation, we used an *in vitro* reconstituted system with synthetic lipid vesicles mimicking the mitochondrial outer membrane (Kuwana et al., 2002; Lovell et al., 2008). While this system lacks certain characteristics of the mitochondrial membrane, most importantly certain relevant protein components (Schafer et al., 2009; Shamas-Din et al., 2013), it has been reliably used to identify properties of the activators and inhibitors of Bax.

To assess the process of Bax activation and membrane insertion we used a single-cysteine mutant of Bax labeled with the dye NBD as a reporter of membrane insertion, as described previously (Kale et al., 2014; Lovell et al., 2008) (Methods). In each set of kinetic experiments, NBD-126C-Bax was incubated with liposomes and an activator (cBid or Bim BH₃ peptide), producing fluorescence timecourses indicating the rate and extent of Bax insertion. To determine the influence of individual components on insertion kinetics, titrations of the system’s three components components were performed (Figure 4.1A). Changes in kinetics were quantified by fitting the NBD-126C-Bax fluorescence timecourses with the two-parameter exponential equation

$$NBD_{norm}(t) = 1 + F_{max}(1 - e^{-kt}) \quad (4.1)$$

with k defining the rate of the insertion process and $1 + F_{max}$ the fold-change NBD-fluorescence at equi-

Mechanism/hypothesis	High [Bax] effect on F_{max}	High [Bax] effect on k
Simple partitioning	No effect	No effect
Bax/membrane binding site	Decrease	Increase
Bax kinetically saturates Bid	No effect	Decrease
Bax oligomerization	Increase	Increase
Observed	Increase	Decrease

Table 4.1: Effects of particular mechanisms on the rate and extent of Bax insertion

librium. Example timecourses with corresponding fits are shown in [Figure 4.1B](#).

4.2.1 WHEN ACTIVATED BY cBid, THE RATE OF BAX MEMBRANE INSERTION SATURATES WITH INCREASED BAX

Surprisingly, when NBD-126C-Bax was titrated at standard concentrations of cBid (~ 20 nM) and lipid vesicles (0.1 mg/mL), we found that the insertion rate k declined with increasing Bax ([Figures 4.1B-D](#)). A previous study using measurements at a 2-hour endpoint showed a similar phenomenon of membrane saturation, for which a stoichiometrically limiting membrane binding site for Bax was proposed as a mechanism ([Satsoura et al., 2012](#)). However, a limitation in Bax binding sites would tend to suppress equilibrium binding of Bax ([Figure 4.1E](#), “Lipo sites”), while the scaling of the fluorescence values with Bax concentration at 1.5, 3 and 6 hours shows that the earlier timepoints show a more significant reduction in Bax binding than the 6 hour endpoint ([Figure 4.1D](#)). In addition, the fitted equilibrium insertion, F_{max} , actually *increased* at the highest concentrations of Bax ([Figure 4.1B](#)). These results suggest that the saturation is primarily kinetic rather than stoichiometric.

To explain the unusual kinetics we observed, we considered several possibilities: first, that high Bax levels were causing the activator cBid to be kinetically saturated, limiting the rate (k) of Bax recruitment ([Figure 4.1E](#), “Bid saturation”); however, this hypothesis could not explain the rise in equilibrium Bax recruitment (F_{max}). On the other hand, we considered the possible role of Bax oligomerization in explaining the increase in equilibrium Bax recruitment ([Figure 4.1E](#), “Bax oligo.”), though this would not account for the reduction in the recruitment rate. As summarized in [Table 4.1](#), none of these mechanisms in isolation made predictions that corresponded with the observed data.




Figure 4.1: Rate and extent of NBD-126C-Bax insertion under unlabeled Bax titration with cBid as the activator.

(A) The *in vitro* model system is comprised of three components: liposomes, Bax, and activator (cBid or Bim BH3). Experiments take the form of titrations of one or more of these components.

(B) Normalized NBD-126C-Bax fluorescence at varying Bax concentrations, using 20 nM cBid and 0.1 mg/mL liposomes. Black lines show fits to Eq. 4.1.

(C) F_{max} and k values resulting from fitting Eq. 4.1 to fluorescence timecourses at each concentration.

(D) Relative fluorescence values at 1.5, 3, and 6 hrs. Plotted values indicate the mean of ten timepoints prior to the specified time-point; error bars indicate the standard deviation over the points.

(E) Predictions of relative Bax insertion at varying Bax concentrations by four mechanistic models.

4.2.2 MEMBRANE ABUNDANCE AFFECTS THE FRACTION OF MEMBRANE-INSERTED BAX, BUT ONLY KINETICALLY

To identify the mechanisms responsible for the kinetics we observed, we performed a series of additional titration experiments intended to isolate the roles of particular constituents in governing the kinetics. We first tested the prior hypothesis that Bax insertion is stoichiometrically limited by a membrane binding site by titrating liposomes while keeping the concentration of NBD-126C-Bax constant (Figure 4.2A). To minimize cBid-dependent effects on kinetics and isolate the Bax-liposome interaction in this experiment we activated Bax using a high concentration ($50\mu\text{M}$) of Bim BH₃ peptide rather than cBid (we revisit the role of cBid below).

The resulting timecourses of relative NBD fluorescence clearly show that the rate of Bax recruitment increases with liposome concentration (Figure 4.2A). We sought to determine the extent to which the liposome concentration affected equilibrium Bax insertion, since the two models of Bax-membrane binding make different predictions: the binding site hypothesis predicts that when binding sites are limited by low liposome concentration, the fraction of Bax bound at equilibrium (F_{max}) will be diminished (Figure 4.2B). On the other hand, the partitioning model predicts that given sufficient time, the amount of Bax recruited at equilibrium will not be dependent on the liposome concentration.

Noting that the fitted F_{max} values at the three highest liposome concentrations were nearly equal, we hypothesized that the maximal fluorescence was independent of the liposome concentration, consistent with the simple partitioning model. Fitting the data with a global F_{max} value across all Bax concentrations yielded good fits, with values for the insertion rate k that increased linearly with liposome concentration, consistent with a pseudo-first-order reaction (Figures 4.2A and 4.2C). This analysis suggests that a mechanism in which Bax is irreversibly recruited to membranes without any stoichiometric limitation is sufficient to explain the major features of the liposome titration data.




Figure 4.2: Rate and extent of NBD-126C-Bax insertion under liposome titration with Bim BH3 peptide as the activator.

(A) Relative NBD-126C-Bax fluorescence at increasing liposome concentrations. Red curves indicate fits to Eq. 4.1.

(B) Predictions of relative Bax insertion under liposome titration by models M1 and M3.

(C) Scaling of fitted k values using a globally fit F_{max} value of 3.85. The green line shows a linear fit to the log-transformed values; the blue line shows a linear fit to the non-transformed values.

(D) Schematic representation of models M1-M4. Models M2 and M4 incorporate the reverse rate k_{r2} , models M1 and M3 do not.

(E) Global fit of model M1 to the liposome titration data.

(F) Relative marginal likelihood values for models M1-M4. Values represent the natural logarithm of the relative probability between each model and the highest probability model M1.

(G) Posterior probability distribution of the number of sites per liposome in model M3 after model calibration.

(H) Posterior probability distribution of the rate constant k_{r2} of model M2 after model calibration.

4.2.3 MECHANISTIC MODELING SHOWS THAT MEMBRANE SITES ARE NOT LIMITING FOR BAX, AND SPONTANEOUS BAX INACTIVATION IS NEGLIGIBLE

Seeking an explanation expressed in terms of elementary reactions, we formulated several mechanistic models for Bax recruitment and examined their fit to the titration data in [Figure 4.2A](#). ([Methods](#)). Prior studies have shown that Bax binding of membranes is a prerequisite for Bax activation in this setting ([Lovell et al., 2008](#)), allowing us to restrict our initial set of models to those of the general form



in which soluble Bax (Bax_s) reversibly binds membranes to form a peripherally membrane-bound form (Bax_m) ([Yethon et al., 2003](#)). This membrane-bound form can either return to solution or undergo an activating conformational change by encountering an activator BH₃-only protein or peptide. To fit the model to the experimental data, we used the relative fluorescence of inserted Bax,

$$F = c_2([Bax_i]/[Bax_o]) \quad (4.3)$$

with $[Bax_o]$ denoting the total concentration of Bax and c_2 the fluorescence of the inserted state relative to aqueous Bax (Bax_s).

The chemical reactions defining each model were formulated as systems of ordinary differential equations by applying the law of mass action and simulating deterministically. For example, the “simple partitioning” model M_I shown in [Figure 4.2D](#), which follows the reaction scheme in [Eq. 4.2](#) but eliminates the reversal of activation by setting $k_{r2} = 0$, yields a system with three rate equations and four parameters, similar to an enzyme-substrate system with the liposomes as the (unsaturable) enzyme ([Chen et al., 2010](#)):

$$d[Bax_s]/dt = -k_{f1}[L][Bax_s] + k_{r1}[Bax_m] \quad (4.4)$$

$$d[Bax_m]/dt = k_{f1}[L][Bax_s] - k_{r1}[Bax_m] - k_{f2}[Bax_m] \quad (4.5)$$

$$d[Bax_i]/dt = k_{f2}[Bax_m] \quad (4.6)$$

Calibration of models to data was performed by Markov chain Monte Carlo sampling, allowing the posterior probability distributions to be estimated for unknown parameters (Gelman et al., 2014). Models were compared by estimating the Bayesian integral known as the marginal likelihood, which specifies the overall probability of a model in light of the data (Calderhead and Girolami, 2009; Girolami et al., 2010). Because parameters and their associated uncertainties are integrated out, this framework allows models with different numbers of unknown parameters to be compared rigorously (Methods).

In the case of model M₁, sample fits from the posterior parameter distribution show that even this simple model fits the liposome titration data quite well (Figure 4.2E). We fit three additional models that incorporated either a non-zero reverse rate k_{r2} (M₂), a limited number of sites per liposome (M₃), or both (M₄) (Figure 4.2D). Models with limited binding sites exhibited substantially worse fits to the data as measured by the marginal likelihoods of the models relative to M₁ (Figure 4.2F). Moreover, the posterior probability density for the number of sites per liposome in M₃ was skewed entirely toward the maximum allowed by the prior (1000 sites per liposome), indicating that it is highly probable that liposome binding sites are not limiting (Figure 4.2G).

To determine whether Bax inactivation or “jumping” (Shamas-Din et al., 2014) was likely to result in diminished Bax recruitment at low liposome concentrations, we examined the posterior distribution of the activation reverse rate k_{r2} in model M₂. The distribution for the reverse rate fell entirely below 10^{-4} sec^{-1} , suggesting that inactivation was not likely to be a factor affecting Bax binding during our experiments (Figure 4.2H). Our upper bound corresponds roughly to the previously published estimate of $2 \times 10^{-4} \text{ sec}^{-1} \pm 1$, which was obtained using a different experimental system (Shamas-Din et al., 2014).

Taken together, these results indicate that the kinetics of Bax activation by a BH₃ peptide under liposome titration can be explained by a very simple partitioning and activation model in which Bax inactivation is negligible and liposome binding sites are not stoichiometrically limiting. Despite not being stoichiometrically limiting, the model indicates that Bax insertion is kinetically very sensitive to liposome concentration, as recruitment is substantially delayed, though not ultimately prevented, at low liposome concentrations.

4.2.4 NBD-126C-BAX FLUORESCENCE INCREASES CONCURRENTLY WITH A REVERSIBLE OLIGOMERIZATION STEP

To further evaluate these models of Bax activation, we performed the converse experiment in which we kept liposomes constant while titrating in unlabeled Bax as a competitor using Bim BH₃ as an activator (Figure 4.3A). For this experiment, the partitioning model M₁ predicts that the insertion kinetics of the labeled Bax will be unaffected by the competitor, while liposome binding site models will show diminished Bax binding as competitor is increased (Figure 4.1E; Table 4.1).

Surprisingly, the results showed that both the rate (k) and extent (F_{max}) of Bax insertion *increase*, rather than decrease, with Bax concentration (Figures 4.3A-B). Though this further supports the conclusion that liposome binding sites are not limiting for Bax, it also indicates that the simple partitioning models M₁ and M₂ do not adequately capture the dynamics of the system.

To explain these results we considered that Bax oligomerization could be responsible for the increase in Bax recruitment at higher Bax concentrations (Figure 4.1E). If the fluorescent state were to increase substantially after oligomerization, higher levels of Bax would shift the equilibrium to favor the highly fluorescent oligomer rather than the less-fluorescent monomer. The addition of Bax oligomerization introduced additional species of Bax into the system requiring us to make further assumptions about which species contributed to the experimentally observed NBD fluorescence, combinatorially expanding the set of hypotheses to include 16 models (Figure 4.3C).

Upon fitting these models to the titration data in Figure 4.3A, we found that while this data alone was not able to discriminate a single best model, it was able to reveal necessary features of the mechanism (Figure 4.3D). As expected, models M₁ and M₂ failed to fit the data because they exhibited no scaling of NBD-126C-Bax fluorescence (Figure 4.1E; Figure 4.3D). Models M₅ and M₇, which respectively contain irreversible or reversible dimerization, also failed in the same fashion: both had an irreversible activation step ($k_{r2} = 0$) so dimerization *after* activation had no effect on the change in NBD fluorescence, which included *both* monomeric and dimeric forms of activated Bax.

Calibration to the liposome titration data (Figure 4.3A) effectively eliminated models M₁, M₂, M₅, M₇, M₉ and M₁₀ from consideration, but the remaining models were separated by smaller differences in their




Figure 4.3: Rate and extent of NBD-126C-Bax insertion under unlabeled Bax titration with Bim BH3 peptide as the activator.

Figure 4.3: (Previous page.) Rate and extent of NBD-126C-Bax insertion under unlabeled Bax titration with Bim BH3 peptide as the activator.

(A) Relative NBD-126C-Bax fluorescence at increasing Bax concentrations. Red curves indicate fits to Eq. 4.1.

(B) F_{max} and k values resulting from fitting Eq. 4.1 to fluorescence timecourses at each concentration.

(C) Schematic representation of models M1-M2 (no dimerization) and M5-M16 (dimerization). Each group of four models (M5-M8, M9-M12, M13-M16) consists of corresponding sets of reactions with either, none, or both of the reverse rates k_{r2} and k_{r3} set to 0. The groups of models differ based on which molecular species contribute to the NBD fluorescence. In models M5-M8 it includes all activated Bax; in M9-M12 only Bax dimers; in M13-M16 monomeric Bax and dimeric Bax contribute separately with distinct fluorescent intensities.

(D) Relative marginal likelihood values for models M1-M2 and M5-M16 after calibration against the Bax titration data in (A). Values represent the natural logarithm of the relative probability between each model and the highest probability model M8. Colors correspond to the model groups in (C).

(E) Relative marginal likelihood values for models M1-M2 and M5-M16 after calibration against the liposome titration data in Figure 4.2. Values represent the natural logarithm of the relative probability between each model and the highest probability model M11. Colors correspond to the model groups in (C).

marginal likelihoods (Figure 4.3D). To further refine this subset we calibrated them independently against the liposome titration data from Figure 4.2 (Figure 4.3E). Here the data discriminated along a different characteristic of the models, showing significantly higher probability for models M9-16 in which the NBD fluorescent species included the dimer. This further eliminated M6 and M8 from consideration.

Taken together, these results show that higher levels of Bax lead to faster and greater fractional recruitment, with the fluorescence increase of NBD-126C-Bax coordinated with dimerization.

4.2.5 THE ACTIVATION OF BAX BY cBid IS NOT “HIT-AND-RUN,” BUT IS SATURABLE WITH AN EFFECTIVE K_M OF ~ 150 nM BAX

Since both titration experiments using a Bim BH₃ peptide to activate Bax showed no evidence of a kinetic saturation attributable to Bax-liposome or Bax-Bax interactions, we considered the possibility that the barrier to Bax recruitment that we observed in our initial experiments (Figure 4.1) was due to features of the cBid/Bax interaction. We therefore performed titrations with unlabeled Bax in the presence of 25 nM NBD-126C-Bax and a range of cBid concentrations (2–80 nM) and fit the resulting kinetic curves with Eq. 4.1 (Figure 4.4A).

For all cBid concentrations, the presence of competing Bax dramatically slowed down the relative NBD-126C-Bax insertion rate leading to apparent saturation (Figure 4.4A). While the predicted F_{max} increased somewhat at high Bax, the rate k diminishes dramatically. To determine if the scaling of the rate k was consistent with a canonical enzymatic role for cBid, we fit each curve individually with a Michaelis-




Figure 4.4: Rate and extent of NBD-126C-Bax insertion under joint titration of cBid and unlabeled Bax.

(A) F_{max} and k values resulting from fitting Eq. 4.1 to fluorescence timecourses at each pair of concentrations.

(B) Insertion rates (k) for 25 nM NBD-126C-Bax and varying cBid concentration (blue line). The red line indicates the expected linear scaling of k based on the lowest two cBid concentrations after assuming a fixed k_{cat} for all cBid concentrations.

(C) FRET between Alexa-568-cBid and DiD-labeled liposomes after 2 hour incubation at 37C, fitted with the lipid binding site and Gouy-Chapman models. Error bars indicate the standard deviation of three replicate wells. Shaded area indicates the 95% confidence interval of model predictions after calibration; the red line indicates the mean prediction.

(D) Predicted cBid binding under liposome titration using the fitted lipid binding site or Gouy-Chapman models. Blue dots indicates the previously published K_D (liposome concentration yielding 50% binding).

(E) Posterior probability distribution for the number of lipid binding sites per liposome after fitting the lipid binding site model.

(F) Posterior probability distribution for parameters of the Gouy-Chapman model after fitting. v indicates the net effective charge on cBid (red line drawn at the calculated value 2); b indicates the ionic strength coefficient (red line at the calculated value 11.5).

Menten equation for the rate of relative product production:

$$V_{norm}([Bax]) = \frac{k_{cat}[cBid]}{K_M + [Bax]} \quad (4.7)$$

The close correspondence of the observed k data with the scaling predicted by Eq. 4.7 suggest that the kinetic barrier to Bax insertion is mediated by an enzymatic mechanism with cBid as the enzyme. However, Eq. 4.7 also predicts that for any concentration of Bax, k will increase linearly with the concentration of cBid, while it clearly does not (Figure 4.4B). Additional cBid does not lead to a proportional increase in Bax recruitment, which explains why the Bax saturation effect was previously observed at a variety of Bax concentrations even though the Bax:cBid ratio was maintained at 5:1 (Satsoura et al., 2012).

4.2.6 BID BINDING TO MEMBRANES IS STOICHIOMETRICALLY LIMITED

We hypothesized that the diminishing effects of cBid were due to a failure of cBid to be stably recruited to membranes at high cBid/liposome ratios. To test this, we measured cBid/membrane binding by FRET using cBid labeled with Alexa 568 as the donor and liposomes labeled with the lipophilic dye DiD as the acceptor (Methods). We incubated 10 nM of the labeled cBid with DiD-liposomes and titrated unlabeled cBid as a competitor (Figure 4.4C). Data from the 2-hour endpoint indicate that unlabeled cBid dramatically reduces the ability of the labeled cBid to bind liposomes. This is in direct contrast with Bax activated by Bim BH₃, where the addition of unlabeled Bax increased, rather than decreased, Bax-membrane association (Figure 4.3A).

The ability of BH₃ proteins to bind membranes is known to be sensitive to the protein and lipid constituents of membranes (Kuwana et al., 2002; Shamas-Din et al., 2015; Shamas-Din et al., 2013), but a stoichiometric limitation on Bid's ability to bind membranes has not been previously described. We considered two possible mechanisms to explain the Bid-membrane binding data we observed. The first is a model positing a limited number of binding sites on the membrane surface, for example a particular lipid species that is required for binding (Z. X. Wang, 1995). The second model is based on the Gouy-Chapman theory describing interactions at charged surfaces and has been previously used to describe the non-ideal membrane binding properties of the pore forming peptide melittin (Schwarz and Beschi-

aschvili, 1989) (Methods). Both of these models produced fairly good fits to the data, though the fit of the Gouy-Chapman model was superior (Figure 4.4C). To determine whether the parameters derived from our fitting procedure were consistent with existing observations of Bid-membrane binding, we simulated a previously published liposome titration experiment in which the fraction of cBid bound was estimated by fluorescence correlation spectroscopy (Shamas-Din et al., 2013) (Figure 4.4D). We found that the lower bound of the 95% confidence interval of our predictions was close to the previously reported K_D of 0.9 nM liposomes for cBid. This correspondence was remarkable given that our FRET experiments were performed in a plate-based assay whereas the published K_D for liposomes was obtained in a cuvette with stirring, which increases association rates and leads to lower apparent K_D values.

We then examined the probability distributions of the fitted parameters of both models to determine if they fell within physically plausible ranges. The binding site model had four free parameters, for the affinity of cBid to the lipid binding site (K), the concentration of lipid binding sites (P_o), the FRET efficiency of the bound state (F_{eff}), and the degree of non-specific binding (N). The posterior distribution of P_o , the concentration of lipid binding sites, showed a range of roughly 1–7 binding sites per liposome, which is implausibly low, given that even the lowest-concentration lipid species, cardiolipin, is present at roughly 8,000 molecules per liposome.

On the other hand, the free parameters of the Gouy-Chapman model fell within plausible ranges. For example, the effective charge on each cBid monomer, denoted v in the model, had values of two or less; the net charge on tBid determined by inspection of the sequence was 2, which in practice represents an upper bound on the actual effective charge of the protein (Roise, 1993). In addition, the dimensionless coefficient b was consistent with values calculated from the ionic strength of the solution. We therefore find the Gouy-Chapman model to be a plausible explanation of the saturation binding we observe for cBid, especially given the important role that electrostatic interactions are already known to play in Bid-membrane interactions.

4.2.7 BAX PORES IN LARGE UNILAMELLAR VESICLES ARE STABLE AND EXHIBIT ALL-OR-NONE RELEASE

We then turned our attention to the kinetics of pore formation, which occurs downstream of Bax insertion (Lovell et al., 2008). An essential first step in characterizing the pore formation mechanism is to determine whether release is *all-or-none*, reflecting stable pores, or *graded*, reflecting transient pores (Figure 4.5A). A prior study using mitochondrial outer membrane vesicles suggested a mechanism at least partially graded (Gillies et al., 2015), whereas a study using giant unilamellar vesicles suggested a stable and growing pore (Bleicken et al., 2013).

To determine the pore formation mechanism in our large unilamellar vesicles we used fluorescence re-quenching, which makes use of the ANTS/DPX dye/quencher pair (Ladokhin et al., 1995). This method exploits the fact that in an all-or-none release mechanism, the degree of quenching of the fluorescent ANTS that remains entrapped, denoted Q_{in} , is invariant to the amount of dye released (Figure 4.5A, blue curve), whereas in a graded release mechanism, ANTS remaining entrapped is gradually dequenched as the overall amount of released dye increases (Figure 4.5A, red curve). In re-quenching analysis, Q_{in} is determined by titrating in additional amounts of the quencher DPX after pore formation has reached steady state: the released ANTS becomes quenched, but entrapped ANTS is not (Ladokhin et al., 1995). Re-quenching analysis has been used to determine the permeabilization mechanism of a number of pore-forming peptides (Wimley et al., 1994; Gregory et al., 2008; Ladokhin et al., 1995).

We performed re-quenching analysis for Bax under a variety of conditions to account for the possibility that the stability of the Bax pore might depend on the strength or type of activation used. Thermal activation of Bax by incubation with liposomes at 43°C yielded a re-quenching curve supporting an all-or-none mechanism, as did the antimicrobial peptide cecropin A, which was previously shown to have an all-or-none mechanism (Gregory et al., 2008) (Figure 4.5B). Activation of Bax by cBid at a variety of different concentrations also indicated an all-or-none mechanism (Figure 4.5C). We therefore concluded that Bax pores in 100 nm large unilamellar vesicles are stable, with all-or-none release.




Figure 4.5: Requenching analysis of Bax pore formation in LUVs.

(A) All-or-none vs. graded permeabilization. All-or-none release (blue curve) results in constant quenching (Q_{in}) across a range of dye release levels (F_{out}); graded release (red curve) results in Q_{in} values that increase with F_{out} .

(B) Q_{in} vs. F_{max} for Bax and cecropin A incubated at 43C.

(C) Q_{in} vs. F_{max} for Bax activated by varying concentrations of cBid.

4.2.8 PERMEABILIZATION OF LIPOSOMES BY BAX SLOWS DOWN WITH TIME AND IS NOT “ENZYMATIC”

We then measured the kinetics of Bax pore formation using ANTS/DPX liposomes as our indicator of liposome permeabilization. As with our Bax insertion studies, we incubated the labeled liposomes with varying concentrations of Bax and an activator BH_3 -only to determine features of the permeabilization mechanism.

In our experiments we noticed that rate of dye release is not constant, but rather slows down with time and plateaus below 100% release (Figure 4.6A-B). This can be seen clearly by plotting the hazard rate $-F'(t)/F(t)$, which measures the instantaneous permeabilization rate over time (Figure 4.6B). A simple enzymatic conversion process is characterized by a hazard rate that is fixed over time, which does not fit the permeabilization kinetic data (Figures 4.6A-B, blue line). On the other hand, the function

$$F(t) = F_{max}(1 - e^{-kt}) \quad (4.8)$$

with a parameter (F_{max}) defining the (submaximal) steady-state permeabilization level, results in very good fits (Figure 4.6A-B, red line). The two-parameter exponential function Eq. 4.8 is similar to those used in previous kinetic studies of Bax-mediated permeabilization (Saito et al., 2000; Schlesinger and Saito, 2006; Kushnareva et al., 2012), indicating that our result is consistent with previous findings. However, the concentration dependence of the extent (rather than the rate) of permeabilization has not previously been analyzed.

Permeabilization curves that reach steady-state levels below 100% release have been observed for pore-forming peptides (Schwarz et al., 1992; Andersson et al., 2007). A graded, transient release mechanism is one possible explanation for this observation, but this was previously ruled out by our requenching experiments (Figure 4.5). An alternative explanation is that stable insertion of Bax into membranes progressively depletes Bax from the solution, slowing release. To test this hypothesis, we pre-incubated cBid and Bax with varying amounts of unlabeled liposomes for 2 hours to allow pore formation to reach steady state, then added in labeled liposomes and measured pore formation for an additional 3 hours (Figure




Figure 4.6: Stable recruitment of Bax to membranes results in submaximal permeabilization.

(A) Representative ANTS/DPX permeabilization curve for 590nM Bax activated by heating to 43°C. Blue curve indicates the best fit to an “enzymatic” (constant rate) permeabilization model. Red curve indicates best fit to the submaximal permeabilization model Eq. 4.8.

(B) Hazard rate of data shown in (A), with corresponding fits.

(C) Permeabilization kinetics following liposome preincubation. During the pre-incubation step, 100 nM cBid, 100 nM Bax, and varying concentrations of unlabeled liposomes were incubated for 2 hours at 37°C; control wells contained only liposomes. After incubation, 1.6 nM ANTS/DPX labeled liposomes were added to the pre-incubation wells; corresponding concentrations of cBid, Bax and ANTS/DPX liposomes were added to control wells. Permeabilization kinetics were then measured by fluorescence spectroscopy for 3 hours. Bar plot shows endpoint permeabilization levels for two replicates at each liposome concentration; inset shows kinetic time-courses for 0.68 nM unlabeled liposomes.

4.6C). Depletion of Bax due to insertion during the pre-incubation phase would result in decreased release in the second phase of the experiment. As a control, we performed experiments in which the labeled and unlabeled liposomes were incubated simultaneously.

When the concentration of liposomes in the pre-incubation phase was small (less than 0.1 nM), the preincubation step had no effect on the extent of release during the measuring step (Figure 4.6C). This is consistent with our findings that the kinetics of Bax recruitment are highly dependent on liposome concentration (Figure 4.2), with low liposome concentrations resulting in incomplete recruitment during the 2 hour incubation. However, intermediate concentrations of unlabeled liposomes almost completely suppressed pore formation, consistent with near-complete recruitment of Bax followed by slow dissociation (Shamas-Din et al., 2014) (Figure 4.2). When concentrations of unlabeled liposomes were substantially higher than the labeled liposomes pore formation was additionally suppressed in the controls due to competition between the two liposome populations for soluble Bax. These results indicate that Bax is substantially depleted from solution after incubation with membranes, which is sufficient to prevent dye release from any vesicles remaining unpermeabilized.

4.2.9 THE DISTRIBUTION OF BAX AMONG LIPOSOMES CANNOT BE INDEPENDENT

The above results indicate that given sufficient time, virtually all soluble Bax will be incorporated into membranes (Figure 4.2). The fraction of permeabilized liposomes at steady state therefore yields a measure of the extent of permeabilization obtained for a known amount of membrane-bound Bax. In principle, this information could be used to estimate the minimal stoichiometry of the pore. However, estimating the “threshold” number of Bax monomers required to form a competent pore not only depends on knowledge of the concentration of membrane-bound Bax, but also the distribution of Bax molecules across the liposome population. The most straightforward assumption is that Bax distributes independently among liposomes, that is, the insertion of Bax in a liposome is independent of whether there is any pre-existing inserted Bax. This type of mechanism gives rise to a Poisson distribution of Bax across liposomes (Figure 4.7A). Liposomes containing greater than the minimal number of Bax molecules for a pore will ultimately be permeabilized, while those below this threshold will remain intact. As the Bax




Figure 4.7: Scaling of maximal permeabilization with Bax/liposome ratio.

(A) The Bax/liposome ratio defines the mean and variance of a Poisson distribution of activated Bax across liposomes. The minimum pore size (vertical line) defines a threshold for permeabilization, shown here at four subunits. As the average Bax/liposome ratio increases, the fraction of permeabilized liposomes increases until reaching 100% (inset plot).

(B) Scaling of predicted maximum dye release (fitted F_{max} values) at varying Bax concentrations after incubation of Bax with 5.2 nM liposomes at 43C for 3 hr (blue line; error bars indicate mean and standard deviation of three experimental replicates). Green line shows the predicted scaling of the maximum release for Poisson-distributed Bax with a minimum pore size of four; red line shows the best Poisson-distribution fit (minimum pore size of 33).

(C) Predicted minimum pore size at each Bax concentration for the data in (B) calculated using the inverse survival function of the Poisson distribution (Methods). The predicted values are fit with a linear equation (red); dashed portion shows extension towards the intercept of 4.2 subunits at [Bax] = 0. Gray dashed line shows the smallest value inferred directly from an experimental measurement (6 subunits).

concentration is increased, the Bax distribution broadens and the mean increases; at high Bax:liposome ratios, all liposomes will be above the threshold, corresponding to full permeabilization (Figure 4.7A, inset).

We performed dye release experiments at a variety of Bax concentrations and estimated the steady-state permeabilization levels by fitting the data with Eq. 4.8. The data show that F_{max} scales slowly with Bax concentration, with only incremental increases in F_{max} for fold-change increases in Bax (Figure 4.7B). Models of pore formation assuming a Poisson distribution of Bax showed much sharper increases in F_{max} and could not fit the data, indicating that the distribution of Bax across liposomes cannot be independent and is influenced by the presence of pre-existing pores.

4.2.10 THE BAX PORE HAS A MINIMUM PORE SIZE OF APPROXIMATELY FOUR SUBUNITS

At low Bax/liposome ratios a soluble Bax molecule is more likely to encounter a liposome with no pores than one with a pre-existing pore, thus making recruitment by auto-activation the less dominant mechanism. Under these conditions the distribution of Bax among liposomes is closer to being independent, allowing inference of the minimal pore size.

We therefore performed a dye release experiment with relatively high liposome concentration (5.2 nM) and Bax titrated over a range of Bax/liposome ratios from approximately 2 to 100. Each dye release curve was fitted with Eq. 4.8 to extract the predicted F_{max} at steady state. As shown in Figure 4.7A, a given Bax/liposome ratio defines a Poisson distribution, with the measured F_{max} values defining the fraction of the distribution below the threshold pore size. The F_{max} values can therefore be used to infer a threshold pore size at each concentration, with the results shown in Figure 4.7C (Methods). Strikingly, the predicted pore sizes fell on a line with higher Bax/liposome ratios resulting in higher predicted pore sizes, an additional indication of the diminishing fractional permeabilization associated with high Bax concentrations. Extrapolation to the low-Bax limit by linear fitting yielded an intercept of 4.2 subunits for the minimum pore size.

The lowest predicted pore size from a measured F_{max} value was 6 subunits, associated with the lowest Bax/liposome ratio of roughly 2 (Figure 4.7C); this value therefore represents an upper bound. Considering these bounds and the evidence that Bax primarily forms oligomers of dimers (Subburaj et al., 2015; Kim et al., 2006), we interpret these results as indicating a minimum pore size of four subunits. Additional experimentation at low Bax concentrations is required to further validate this estimate and to clarify the types of auto-activation mechanisms that produce the scaling curves seen in Figure 4.7C.

4.3 DISCUSSION

Many aspects of the mechanism of apoptotic pore formation mechanism have been made clear in recent years: an activator BH3-only protein, such as truncated Bid or Bim, is generally required to trigger a series of conformational changes in either Bax or Bak that allow it to oligomerize and form pores in the membrane. Recent structural studies of Bax and Bak oligomers have provided evidence of the conformations of these pore-forming molecules in the membrane (Czabotar et al., 2013; Bleicken et al., 2014). However, despite progress in determining the structure of Bcl-2 proteins and their network of interactions, there remain a number of unresolved questions. In addition to well-recognized gaps in the known mechanism, there are also inconsistencies and ambiguities that manifest when integrating information from disparate studies.

In this study we use data from *in vitro* experiments to develop dynamic models of the apoptotic pore formation process that refine and integrate knowledge about mechanisms. We characterize the roles of lipid, Bax and Bid abundance in controlling the rate and extent of Bax insertion, identifying a previously unknown stoichiometric limitation on the ability of cBid to bind membranes, which we attribute to electrostatic interference among Bid monomers. Dye release experiments indicate that the Bax pore is stable and grows by recruiting Bax monomers but has a minimum size of roughly four subunits.

Bax auto-activation has been proposed to play a major role in the “all-or-none” nature of mitochondrial permeabilization (Cui et al., 2008). In this interpretation, Bax auto-activation is responsible for propagating an initial pore formation event into widespread permeabilization across the entire cell. While the initial propensity for permeabilization would be low, it would accelerate as positive feedback triggered further increases in Bax recruitment and permeabilization. This interpretation is not supported by our data, which shows that the rate of Bax recruitment to membranes is well-fit by a single exponential function, indicating a constant recruitment rate (Figure 4.2). In addition, the rate of permeabilization does not accelerate but instead *decreases* with time (Figure 4.6). At the same time, we show that pore formation is not independent, suggesting a role for auto-activation in favoring recruitment of Bax to existing pores (Figure 4.7). We therefore conclude that the primary role played by Bax auto-activation is the *local* expansion of pores rather than the dissemination of activated Bax to other mitochondrial sites. As our data show, this mechanism *reduces* the extent of permeabilization in the population of vesicles due to the highly stable recruitment of Bax to membranes.

One difficulty in applying the knowledge gained from quantitative studies of Bax pore formation *in vitro* to interpreting data from cells has been the apparent difference in timescales between the two settings: Bax pore formation *in vitro* occurs over a span of 1-3 hours (e.g., Figure 4.6), whereas in cells, widespread permeabilization of mitochondria occurs in minutes (Goldstein et al., 2000; Rehm et al., 2002; Albeck et al., 2008a). Protein constituents of the mitochondrial outer membrane that are lacking in the *in vitro* setting are one potential cause of this discrepancy (Kushnareva et al., 2012; Schafer et al., 2009). While mitochondrial proteins undoubtedly play a role, it is noteworthy that permeabilization of mitochondrial outer membrane vesicles and depolarization of mitochondria by BH₃ peptides also take

far longer than permeabilization in cells (Kushnareva et al., 2012; Sarosiek et al., 2013).

We speculate that a major cause of the difference in timescales is the different relationship between lipid abundance and compartment size between the two settings. In the *in vitro* setting (whether synthetic liposomes, outer membrane vesicles, or isolated mitochondria), the vesicle compartments tend to be smaller than the (possibly highly interconnected) mitochondrial compartments of cells. Moreover, an increase in the amount of vesicles *in vitro* not only increases the recruitment of Bid and Bax to membranes (thereby accelerating pore formation, Figure 4.2) but also increases the number of compartments, which decreases the overall (proportional) permeabilization. In the cell, on the other hand, mitochondria occupy roughly 10% of the cell volume but number only in the hundreds, with the number changing dynamically due to fission and fusion (Posakony et al., 1977). The cell is therefore a high lipid, low compartment number environment, in which Bid and Bax recruitment is favorable and fractional permeabilization occurs rapidly. In this interpretation, the “all-or-none” nature of pore formation is attributable not to auto-activation, but instead to “pore overshoot” resulting from the rapid recruitment of far more Bax than required to permeabilize all mitochondrial compartments (Dussman et al., 2009; Albeck et al., 2008b). In future work, we aim to determine whether our models, after calibrating to *in vitro* data, can quantitatively predict rates of permeabilization in cells after accounting for these differences in mitochondrial size and concentration.

It is currently unknown whether our finding that Bid recruitment to membranes is limited at high Bid:membrane ratios also holds for binding to native mitochondrial membranes, or for other activators such as Bim. However, electrostatic interactions with negatively charged headgroups have been shown to play a role in determining the affinity of both Bim and cBid for membranes, suggesting that unfavorable interactions at surfaces could also affect both proteins (Shamas-Din et al., 2015). The degree of electrostatic interference is likely to be sensitive to the size (and hence the curvature) of the vesicles used, with smaller vesicles resulting in greater electrostatic interference among proteins, allowing us to test this hypothesis further. Addressing these questions is a focus of future work.

In fitting our mechanistic models to our Bax titration data, we found that mechanisms tended to be corroborated across datasets. For example, our Bax titration data revealed that NBD-126C-Bax insertion was coincident with dimerization (Figure 4.3D). The same dimerization models also showed dramatically

improved fit for our initial liposome titration dataset even though simpler binding models M₁ and M₂ had already shown very good fits (Figure 4.3E). This shows that kinetic data of this type tends to contain considerably more information about mechanisms than can be extracted by qualitative interpretation. We expect that this type of analytical approach will become increasingly necessary for interpreting large datasets involving complex cellular perturbations (e.g., Vidović et al., 2014). Even in our relatively simple *in vitro* setting, we find that the context—concentrations of various constituents—dramatically affects the mechanisms that predominate, with each experiment therefore yielding only a partial view of the overall mechanism. Systematic evaluation of models against complex datasets helps to assemble these partial views into a coherent mechanistic picture.

4.4 METHODS

4.4.1 PROTEIN PURIFICATION AND FLUORESCENCE SPECTROSCOPY

Procedures for purifying and labeling cBid and Bax, preparing lipid vesicles, and measuring NBD and ANTS by fluorescence spectroscopy were described in detail previously (Kale et al., 2014; Lovell et al., 2008). Kinetic measurements were made on a FlexStation III (Molecular Devices).

4.4.2 MODELING AND DATA ANALYSIS

Raw data and all code are available in an open-source format on GitHub at <https://github.com/johnbachman/tBidBaxLipo>. Mechanistic models were specified using PySB (Lopez et al., 2013). Numerical simulation and parameter estimation was performed using Markov chain Monte Carlo as described in Chapter 3. Broad prior distributions were used based on the types of rate constants and previously published estimates.

Experimental error for each curve in the titrations was estimated by fitting a polynomial function to the latter part of the curve as it approached equilibrium and then calculating the standard error of the residuals after checking that the residuals followed a Gaussian distribution.

4.4.3 MODELS OF BID-MEMBRANE BINDING

For the lipid binding site model of Bid-membrane binding, we used an exact solution for competitive binding between two ligands to a target (Z. X. Wang, 1995). In this model, the target “receptor” molecule is denoted P , B is the labeled ligand, and A is the titrated competitor. We assume that the addition of the label does not affect binding, that is, that the affinity of A and B for P is equal. We also assume that the spectroscopic signal of the labeled ligand B (Alexa 568 cBid in this case) is unaffected by binding to P .

The mathematical expression for the concentration of the bound, labeled complex $[PB]$ in terms of the ligand-target affinity K and the initial concentrations A_o , B_o , and P_o is given by equation 15 of Wang (Z. X. Wang, 1995). The predicted FRET signal is then

$$\text{FRET} = \frac{F_{eff}}{B_o} ([PB] + NB_o) \quad (4.9)$$

with F_{eff} denoting the FRET efficiency associated with the bound complex PB and N the degree of non-specific binding of B . The resulting model has three parameters: P_o , the number of (lipid) binding sites, F_{eff} and N .

The Gouy-Chapman model of membrane binding was drawn from Schwarz and Beschiaschvili, 1989. In this model, the partitioning of cBid between solution and membranes is described by

$$r = (\Gamma/\alpha) \cdot [cBid_o] \quad (4.10)$$

where r is the concentration of membrane-bound cBid, $[cBid_o]$ is the total solute (cBid) concentration, Γ is the (concentration-independent) partitioning coefficient and α is an activity coefficient describing unfavorable interactions among the solute molecules.

Application of the Gouy-Chapman theory gives an expression for the activity coefficient α ,

$$\ln \alpha = 2v \cdot \sinh^{-1}(vbr) \quad (4.11)$$

in which v is the effective number of charges per cBid monomer and b is a dimensionless coefficient de-

pendent primarily on the ionic strength of the solution (Schwarz and Beschiaschvili, 1989).

4.4.4 CALCULATION OF MINIMUM PORE SIZES

A Poisson distribution is defined by the parameter λ which specifies the mean number of “instances” (in this case, activated Bax molecules) observed within a given “interval” (in this case, a single liposome). The probability of observing k instances within an interval (i.e., a liposome with k Bax molecules) is given by the probability mass function:

$$\frac{\lambda^k}{k!} e^{-\lambda} \quad (4.12)$$

For a given λ , the cumulative distribution function (CDF) of the Poisson distribution gives $P(X \leq k)$, the probability of observing k or fewer instances within an interval—here, the probability of a liposome with k or fewer Bax molecules. If we define $r = k + 1$ as our threshold pore size, the CDF gives the fraction of liposomes expected to remain unpermeabilized.

The *survival function*, $SF(k; \lambda) = P(X > k) = 1 - P(X \leq k)$, therefore gives the probability of a liposome having greater than k Bax molecules. For a given minimum pore size r and Bax/liposome ratio λ , the survival function $SF(r - 1; \lambda)$ therefore returns the expected dye release. The *inverse survival function* $ISF(P; \lambda)$ returns the value of k corresponding to a given λ and fraction permeabilized P . In this analysis we use the implementation of the Poisson ISF in the Python package `scipy`, `scipy.stats.poisson.isf` (Oliphant, 2007).

References

- Albeck, J. G., J. M. Burke, B. B. Aldridge, M. Zhang, D. A. Lauffenburger, and P. K. Sorger (2008a). “Quantitative analysis of pathways controlling extrinsic apoptosis in single cells”. *Molecular Cell* 30.1 (Apr. 2008), pp. 11–25.
- Albeck, J. G., J. M. Burke, S. L. Spencer, D. A. Lauffenburger, and P. K. Sorger (2008b). “Modeling a snap-action, variable-delay switch controlling extrinsic cell death”. *PLoS Biology* 6.12 (Dec. 2008), pp. 2831–2852.
- Andersson, A., J. Danielsson, A. Gräslund, and L. Måler (2007). “Kinetic models for peptide-induced leakage from vesicles and cells.” *European biophysics journal : EBJ* 36.6 (July 2007), pp. 621–635.
- Billen, L. P., C. L. Kokoski, J. F. Lovell, B. Leber, and D. W. Andrews (2008). “Bcl-XL inhibits membrane permeabilization by competing with Bax”. *PLoS Biology* 6.6 (June 2008), e147.
- Bleicken, S., G. Jeschke, C. Stegmüller, R. Salvador-Gallego, A. J. García-Sáez, and E. Bordinon (2014). “Structural model of active Bax at the membrane.” *Molecular Cell* 56.4 (Nov. 2014), pp. 496–505.
- Bleicken, S., O. Landeta, A. Landajuela, G. Basañez, and A. J. García-Sáez (2013). “Proapoptotic Bax and Bak proteins form stable protein-permeable pores of tunable size.” *Journal of Biological Chemistry* 288.46 (Nov. 2013), pp. 33241–33252.
- Calderhead, B. and M. Girolami (2009). “Estimating Bayes factors via thermodynamic integration and population MCMC”. *Computational Statistics & Data Analysis* 53.12 (Oct. 2009), pp. 4028–4045.
- Chen, W. W., M. Niepel, and P. K. Sorger (2010). “Classic and contemporary approaches to modeling biochemical reactions”. *Genes & Development* 24.17 (Sept. 2010), pp. 1861–1875.
- Chonghaile, T. N. et al. (2011). “Pretreatment Mitochondrial Priming Correlates with Clinical Response to Cytotoxic Chemotherapy”. *Science (New York, NY)* 334.6059 (Nov. 2011), pp. 1129–1133.
- Cui, J., C. Chen, H. Lu, T. Sun, and P. Shen (2008). “Two independent positive feedbacks and bistability in the Bcl-2 apoptotic switch”. *PLoS ONE* 3.1, e1469.
- Czabotar, P. E. et al. (2013). “Bax crystal structures reveal how BH3 domains activate Bax and nucleate its oligomerization to induce apoptosis.” *Cell* 152.3 (Jan. 2013), pp. 519–531.
- Dai, H., A. Smith, X. W. Meng, P. A. Schneider, Y.-P. Pang, and S. H. Kaufmann (2011). “Transient binding of an activator BH3 domain to the Bak BH3-binding groove initiates Bak oligomerization.” *The Journal of Cell Biology* 194.1 (July 2011), pp. 39–48.
- Dewson, G., S. Ma, P. Frederick, C. Hockings, I. Tan, T. Kratina, and R. M. Kluck (2012). “Bax dimerizes via a symmetric BH3:groove interface during apoptosis.” *Cell Death and Differentiation* 19.4 (Apr. 2012), pp. 661–670.

- Dussman, H., M. Rehm, C. G. Concannon, S. Anguissola, M. Wurstle, S. Kacmar, P. Voiler, H. J. Huber, and J. H. M. Prehn (2009). "Single-cell quantification of Bax activation and mathematical modelling suggest pore formation on minimal mitochondrial Bax accumulation". *Cell Death and Differentiation* (Nov. 2009), pp. 1–13.
- Edlich, F., S. Banerjee, M. Suzuki, M. M. Cleland, D. Arnoult, C. Wang, A. Neutzner, N. Tjandra, and R. J. Youle (2011). "Bcl-x(L) retrotranslocates Bax from the mitochondria into the cytosol." *Cell* 145.1 (Apr. 2011), pp. 104–116.
- Gavathiotis, E., D. E. Reyna, J. A. Bellairs, E. S. Leshchiner, and L. D. Walensky (2012). "Direct and selective small-molecule activation of proapoptotic BAX." *Nature Chemical Biology* 8.7 (July 2012), pp. 639–645.
- Gavathiotis, E., D. E. Reyna, M. L. Davis, G. H. Bird, and L. D. Walensky (2010). "BH₃-Triggered Structural Reorganization Drives the Activation of Proapoptotic BAX". *Molecular Cell* 40.3 (Nov. 2010), pp. 481–492.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (2014). *Bayesian data analysis*. Chapman and Hall/CRC.
- George, N. M., J. J. D. Evans, and X. Luo (2007). "A three-helix homo-oligomerization domain containing BH₃ and BH₁ is responsible for the apoptotic activity of Bax". *Genes & Development* 21.15 (Aug. 2007), pp. 1937–1948.
- Gillies, L. A., H. Du, B. Peters, C. M. Knudson, D. D. Newmeyer, and T. Kuwana (2015). "Visual and functional demonstration of growing Bax-induced pores in mitochondrial outer membranes". *Molecular Biology of the Cell* 26.2 (Jan. 2015), pp. 339–349.
- Girolami, M., B. Calderhead, and V. Vyshemirsky (2010). "System Identification and Model Ranking: The Bayesian Perspective Learning and Inference". In: *Learning and Inference in Computational Systems Biology*. MIT Press.
- Goldstein, J. C., N. J. Waterhouse, P. Juin, G. I. Evan, and D. R. Green (2000). "The coordinate release of cytochrome c during apoptosis is rapid, complete and kinetically invariant". *Nature Cell Biology* 2.3 (Mar. 2000), pp. 156–162.
- Gregory, S. M., A. Cavenaugh, V. Journigan, A. Pokorny, and P. F. F. Almeida (2008). "A quantitative model for the all-or-none permeabilization of phospholipid vesicles by the antimicrobial peptide cecropin A". *Biophysical Journal* 94.5 (Mar. 2008), pp. 1667–1680.
- Kale, J., X. Chi, B. Leber, and D. Andrews (2014). "Examining the molecular mechanism of bcl-2 family proteins at membranes by fluorescence spectroscopy." *Methods in enzymology* 544, pp. 1–23.
- Kim, H., M. Rafiuddin-Shah, H.-C. Tu, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2006). "Hierarchical regulation of mitochondrion-dependent apoptosis by BCL-2 subfamilies". *Nature Cell Biology* 8.12 (Dec. 2006), pp. 1348–1358.

- Kim, H., H.-C. Tu, D. Ren, O. Takeuchi, J. R. Jeffers, G. P. Zambetti, J. J.-D. Hsieh, and E. H.-Y. Cheng (2009). "Stepwise activation of BAX and BAK by tBID, BIM, and PUMA initiates mitochondrial apoptosis". *Molecular Cell* 36.3 (Nov. 2009), pp. 487–499.
- Kushnareva, Y., A. Y. Andreyev, T. Kuwana, and D. D. Newmeyer (2012). "Bax activation initiates the assembly of a multimeric catalyst that facilitates Bax pore formation in mitochondrial outer membranes." *PLoS Biology* 10.9, e1001394.
- Kuwana, T., M. R. Mackey, G. Perkins, M. H. Ellisman, M. Latterich, R. Schneider, D. R. Green, and D. D. Newmeyer (2002). "Bid, Bax, and lipids cooperate to form supramolecular openings in the outer mitochondrial membrane." *Cell* 111.3 (Nov. 2002), pp. 331–342.
- Ladokhin, A. S., W. C. Wimley, and S. H. White (1995). "Leakage of membrane vesicle contents: determination of mechanism using fluorescence reequenching". *Biophysical Journal* 69.5 (Nov. 1995), pp. 1964–1971.
- Leverson, J. D. et al. (2015). "Potent and selective small-molecule MCL-1 inhibitors demonstrate on-target cancer cell killing activity as single agents and in combination with ABT-263 (navitoclax)." *Cell death & disease* 6, e1590.
- Lopez, C. F., J. L. Muhlich, J. A. Bachman, and P. K. Sorger (2013). "Programming biological models in Python using PySB." *Molecular Systems Biology* 9, p. 646.
- Lovell, J. F., L. P. Billen, S. Bindner, A. Shamas-Din, C. Fradin, B. Leber, and D. W. Andrews (2008). "Membrane binding by tBid initiates an ordered series of events culminating in membrane permeabilization by Bax". *Cell* 135.6 (Dec. 2008), pp. 1074–1084.
- Martinez-Caballero, S., L. M. Dejean, M. S. Kinnally, K. J. Oh, C. A. Mannella, and K. W. Kinnally (2009). "Assembly of the Mitochondrial Apoptosis-induced Channel, MAC". *The Journal of biological chemistry* 284.18 (May 2009), pp. 12235–12245.
- Nechushtan, A., C. L. Smith, I. Lamensdorf, S. H. Yoon, and R. J. Youle (2001). "Bax and Bak coalesce into novel mitochondria-associated clusters during apoptosis." *Journal of Cell Biology* 153.6 (June 2001), pp. 1265–1276.
- Oliphant, T. E. (2007). "Python for Scientific Computing". *Computing in Science & Engineering* 9.3, pp. 10–20.
- Oltersdorf, T. et al. (2005). "An inhibitor of Bcl-2 family proteins induces regression of solid tumours". *Nature* 435.7042 (June 2005), pp. 677–681.
- Posakony, J. W., J. M. England, and G. Attardi (1977). "Mitochondrial growth and division during the cell cycle in HeLa cells." *The Journal of Cell Biology* 74.2 (Aug. 1977), pp. 468–491.
- Rehm, M., H. Dussmann, R. U. Janicke, J. M. Tavaré, D. Kogel, and J. H. M. Prehn (2002). "Single-cell fluorescence resonance energy transfer analysis demonstrates that caspase activation during apoptosis is a rapid process. Role of caspase-3". *The Journal of biological chemistry* 277.27 (July 2002), pp. 24506–24514.

- Roise, D. (1993). "Incorporation of Peptides into Membranes". In: *Thermodynamics of Membrane Receptors and Channels*. Ed. by M. B. Jackson. Boca Raton: CRC Press.
- Saito, M., S. J. Korsmeyer, and P. H. Schlesinger (2000). "BAX-dependent transport of cytochrome c reconstituted in pure liposomes." *Nature Cell Biology* 2.8 (Aug. 2000), pp. 553–555.
- Sarosiek, K. A. et al. (2013). "BID preferentially activates BAK while BIM preferentially activates BAX, affecting chemotherapy response." *Molecular Cell* 51.6 (Sept. 2013), pp. 751–765.
- Satsoura, D., N. Kučerka, S. Shivakumar, J. Pencer, C. Griffiths, B. Leber, D. W. Andrews, J. Katsaras, and C. Fradin (2012). "Interaction of the full-length Bax protein with biomimetic mitochondrial liposomes: a small-angle neutron scattering and fluorescence study." *Biochimica et biophysica acta* 1818.3 (Mar. 2012), pp. 384–401.
- Schafer, B., J. Quispe, V. Choudhary, J. E. Chipuk, T. G. Ajero, H. Du, R. Schneider, and T. Kuwana (2009). "Mitochondrial outer membrane proteins assist Bid in Bax-mediated lipidic pore formation." *Molecular Biology of the Cell* 20.8 (Apr. 2009), pp. 2276–2285.
- Schlesinger, P. H. and M. Saito (2006). "The Bax pore in liposomes, Biophysics". *Cell Death and Differentiation* 13.8 (Aug. 2006), pp. 1403–1408.
- Schwarz, G. and G. Beschiaschvili (1989). "Thermodynamic and kinetic studies on the association of melittin with a phospholipid bilayer". *Biochimica et biophysica acta* 979.1 (Feb. 1989), pp. 82–90.
- Schwarz, G., R. T. Zong, and T. Popescu (1992). "Kinetics of melittin induced pore formation in the membrane of lipid vesicles". *Biochimica et biophysica acta* 1110.1 (Sept. 1992), pp. 97–104.
- Shamas-Din, A., D. Satsoura, O. Khan, W. Zhu, B. Leber, C. Fradin, and D. W. Andrews (2014). "Multiple partners can kiss-and-run: Bax transfers between multiple membranes and permeabilizes those primed by tBid." *Cell death & disease* 5, e1277.
- Shamas-Din, A., S. Bindner, X. Chi, B. Leber, D. W. Andrews, and C. Fradin (2015). "Distinct lipid effects on tBid and Bim activation of membrane permeabilization by pro-apoptotic Bax." *The Biochemical journal* 467.3 (May 2015), pp. 495–505.
- Shamas-Din, A., S. Bindner, W. Zhu, Y. Zaltsman, C. Campbell, A. Gross, B. Leber, D. W. Andrews, and C. Fradin (2013). "tBid undergoes multiple conformational changes at the membrane required for Bax activation." *Journal of Biological Chemistry* 288.30 (July 2013), pp. 22111–22127.
- Souers, A. J. et al. (2013). "ABT-199, a potent and selective BCL-2 inhibitor, achieves antitumor activity while sparing platelets." *Nature medicine* 19.2 (Feb. 2013), pp. 202–208.
- Subburaj, Y., K. Cosentino, M. Axmann, E. Pedrueza-Villalmanzo, E. Hermann, S. Bleicken, J. Spatz, and A. J. García-Sáez (2015). "Bax monomers form dimer units in the membrane that further self-assemble into multiple oligomeric species". *Nature Communications* 6 (Aug. 2015).
- Tait, S. W. G. and D. R. Green (2010). "Mitochondria and cell death: outer membrane permeabilization and beyond". *Nature Reviews Molecular Cell Biology* 11.9 (Sept. 2010), pp. 621–632.

- Tan, C., P. J. Dlugosz, J. Peng, Z. Zhang, S. M. Lapolla, S. M. Plafker, D. W. Andrews, and J. Lin (2006). "Auto-activation of the apoptosis protein Bax increases mitochondrial membrane permeability and is inhibited by Bcl-2". *The Journal of biological chemistry* 281.21 (May 2006), pp. 14764–14775.
- Todt, F., Z. Cakir, F. Reichenbach, R. J. Youle, and F. Edlich (2013). "The C-terminal helix of Bcl-x(L) mediates Bax retrotranslocation from the mitochondria." *Cell Death and Differentiation* 20.2 (Feb. 2013), pp. 333–342.
- Vidović, D., A. Koleti, and S. C. Schürer (2014). "Large-scale integration of small molecule-induced genome-wide transcriptional responses, Kinome-wide binding affinities and cell-growth inhibition profiles reveal global trends characterizing systems-level drug action." *Frontiers in genetics* 5, p. 342.
- Vo, T.-T., J. Ryan, R. Carrasco, D. Neuberg, D. J. Rossi, R. M. Stone, D. J. Deangelo, M. G. Frattini, and A. Letai (2012). "Relative mitochondrial priming of myeloblasts and normal HSCs determines chemotherapeutic success in AML." *Cell* 151.2 (Oct. 2012), pp. 344–355.
- Volkman, N., F. M. Marassi, D. D. Newmeyer, and D. Hanein (2013). "The rheostat in the membrane: BCL-2 family proteins and apoptosis." *Cell Death and Differentiation* (Oct. 2013).
- Wang, K., X. M. Yin, D. T. Chao, C. L. Milliman, and S. J. Korsmeyer (1996). "BID: a novel BH₃ domain-only death agonist." *Genes & Development* 10.22 (Nov. 1996), pp. 2859–2869.
- Wang, Z. X. (1995). "An exact mathematical expression for describing competitive binding of two different ligands to a protein molecule." *FEBS letters* 360.2 (Feb. 1995), pp. 111–114.
- Wimley, W. C., M. E. Selsted, and S. H. White (1994). "Interactions between human defensins and lipid bilayers: evidence for formation of multimeric pores." *Protein science : a publication of the Protein Society* 3.9 (Sept. 1994), pp. 1362–1373.
- Yethon, J. A., R. F. Epand, B. Leber, R. M. Epand, and D. W. Andrews (2003). "Interaction with a membrane surface triggers a reversible conformational change in Bax normally associated with induction of apoptosis." *The Journal of biological chemistry* 278.49 (Dec. 2003), pp. 48935–48941.

5

Conclusion

In this thesis I describe a computational framework for efficiently enumerating alternative biochemical hypotheses and apply it to analyzing the dynamics of apoptotic pore formation. Use of an *in vitro* experimental system allows systematic experimentation involving many different Bax mutants and wide variation of the concentrations of components. Collections of programmatically enumerated alternative models are calibrated against these complex datasets, identifying and clarifying pore formation mechanisms.

The *in vitro* experiments in this thesis highlight the degree to which mechanisms governing dynamics of biochemical systems are context dependent even in highly simplified settings. For example, the Bax insertion kinetics in [Chapter 4](#) were found to depend on Bax/membrane, Bid/membrane, Bid/Bax, and Bax/Bax interactions, with the relative contributions of these mechanisms depending on precise concentrations. At the same time, we showed that these influences on the dynamics could be accounted for by assembling and refining an ensemble of models.

There are two practical challenges in making these approaches scale to account for the more extensive

and challenging context-dependence in cellular settings. The first is to develop ways to make models more trustworthy; the second is to make the model building process more efficient.

5.1 MAKING MODELS MORE TRUSTWORTHY

Trustworthiness, for a model, refers to the confidence a human user has in using the model in a particular context. This depends on multiple factors, including knowing answers to the following questions:

1. What elements are in the model?
2. Why does the model include the elements that it does?
3. What is the range of phenomena covered by the model?

The first of these questions deals with model *transparency*. In [Chapter 2](#) I discussed how conceptual approaches from programming can make models more transparent and thus more trustworthy and reusable. The second question, about *why* the model contains what it does, demands an understanding of the evidentiary basis behind the mechanistic assertions in the model. This depends on a detailed accounting for the empirical evidence supporting the facts in the model, grounded in the primary scientific literature.

Drawing again from approaches to computer programming, I propose that this requirement for model trustworthiness be met by adopting *literate programming*, a style of programming in which the textual description of the program's elements, written for human readers, is the primary organizing feature, with the program code embedded within it ([Knuth, 1984](#)). A “literate model” for biology would be similar to an extensive review of biological system, with verbal descriptions of mechanisms supported by references to experimental evidence. These mechanisms would then be encoded into one or more representations in a formal modeling language, allowing the human reader to compare the formal and informal representations side-by-side. One example of this type of approach is the Yeast Pheromone Model, developed by Ty Thomson (<http://yeastpheromonemodel.org>). We are currently developing this modeling approach in the context of an extensive model of the Ras signaling pathway (<http://rasmodel.org>).

The third criterion for model trustworthiness is concerned with validating the model and establishing its explanatory scope. In the traditional sense, validation involves fitting a model to data to see if (or how well) it fits. However, as models grow more complex, it becomes important to know not only whether the model explains the dataset currently being considered, but also if the model is consistent with prior observations from the literature. This involves setting up *in silico* simulations of previous experiments to see if the model can retrodict their findings.

The purpose of this more extended validation is two-fold: first, to determine whether the conclusions of a model are consistent with prior findings. If so, this increases confidence in the model's ability to generalize, making its future predictions more trustworthy. For example, in [Chapter 4](#), I showed that cBid/liposome binding models calibrated to data from a cBid titration experiment made predictions for a liposome titration experiment consistent with previously published results.

The second goal of doing this type of validation is to establish the explanatory scope of a model—evaluating models against a wide variety of phenomena reveals the types of questions that they can be used to answer. In [Chapter 4](#) I showed that the simplest activation model (M_I) was effective at predicting Bax recruitment for a variety of liposome concentrations, but not for a wide variety of Bax concentrations or for activators other than the Bim BH₃ peptide. Empirically establishing the scope of the simpler model allows it to be used with confidence in the appropriate context.

Ultimately, models could be subject to an ongoing, semi-automated validation process, encompassing both new and previous data. In an analogous matter to what is known as *regression testing* in software, a suite of tests could reveal not only where models fail to meet the current specifications, but also identify when the addition of features have decreased performance on prior datasets.

Extensive documentation and ongoing validation would allow models to be brought to bear more confidently on a wide variety of complex problems in biology. Gaps in models could be used to efficiently identify new hypotheses, and model predictions could be used to identify new phenomena and design effective experiments, accelerating the mechanistic discovery process.

5.2 ENLISTING MACHINES FOR EFFICIENT MODELING

Biological modeling is a very time and labor-intensive process. Modeling requires extensive knowledge not only in a particular biological subfield but also of modeling tools and formalisms. In the pre-systems biology era, there was a saying that encapsulated the reductionist approach: “one gene, one Ph.D.” In the current state of modeling in systems biology, an equivalent observation would be “one model, one Ph.D.”

There has been increasing interest in using machines to build models automatically, either from the scientific literature or from pathway databases (Büchel et al., 2013; Cohen, 2015). Many tasks in modeling are routine and repetitive, especially as they relate to encoding names and features of genes and proteins. Modeling software could be developed to automatically assemble, for any given gene, an agent structure suitable for rule-based modeling, including relevant active sites, interaction domains and post-translational modification sites. Machines could also be enlisted to search the literature for the evidence behind interactions chosen by the modeler (Kemper et al., 2010). Machine assistance would help reduce unintentional bias in model scope arising mainly out of human modelers’ necessarily limited domain knowledge. “Recommender” software could make intelligent suggestions about the genes, proteins, and interactions that should be considered for a model of a particular phenomenon. The semi-automated model validation described above could itself be considered a form of machine-assisted modeling.

Current approaches to machine-assisted modeling function largely as one-way pipelines, processing large corpora to extract interactions or assemble models (Büchel et al., 2013). A complementary approach would be for the biologist to engage with an intelligent agent in a dialog, interacting collaboratively to construct and refine models of biological processes (Carvunis and Ideker, 2014). Such an approach could be very powerful as it would allow the software agents to refine the results of their analyses based on the role of the human user and the goals of the interaction, rather than simply returning a large “hairball” for the human to prioritize.

The prospect of using sophisticated software to build models raises questions about what the actual model artifact would be in this scenario. In the programmatic approach to modeling described in Chapter 2 the model artifact is a computer program, which contains within it all of the mechanistic elements

deemed relevant. In a machine-modeling scenario, the model artifact results from human guidance to a software pipeline that extracts and assembles mechanistic elements into a self-consistent model. The guidance given to the machine is analogous to a search query, and the model that is produced to a complex set of results. For a human user, the search term is too vague to comprehend the resulting model, while on the other hand the machine-assembled result may be too detailed and complex.

By leveraging sophisticated software to help build and curate models, mechanistic explanations of biological processes may be able to scale to meet the growing demands of explaining larger and more complex datasets. Scientists could efficiently explore mechanistic explanations of data formally without having to have deep expertise in neighboring subfields of biology or modeling methods. If successful, such tools could streamline both basic and applied research in biomedicine and serve as the foundation for applications in personalized medicine.

References

- Büchel, F. et al. (2013). “Path2Models: large-scale generation of computational models from biochemical pathway maps.” *BMC systems biology* 7, p. 116.
- Carvunis, A.-R. and T. Ideker (2014). “Siri of the Cell: What Biology Could Learn from the iPhone”. *Cell* 157.3 (Apr. 2014), pp. 534–538.
- Cohen, P. R. (2015). “DARPA’s Big Mechanism program.” *Physical biology* 12.4 (July 2015), p. 045008.
- Kemper, B., T. Matsuzaki, Y. Matsuoka, Y. Tsuruoka, H. Kitano, S. Ananiadou, and J. Tsujii (2010). “PathText: a text mining integrator for biological pathway visualizations.” *Bioinformatics (Oxford, England)* 26.12 (June 2010), pp. 1374–81.
- Knuth, D. E. (1984). “Literate Programming”. *The Computer Journal* 27.2 (Jan. 1984), pp. 97–111.



Optimizing ring assembly reveals the strength of weak interactions

AUTHORSHIP CONTRIBUTION

I chose the focus of the study (ring-like complexes), formulated the mathematical models, and ran kinetic simulations. I identified the phenomenon of the kinetic plateau and collaborated closely with E. Deeds on analytical and numerical approaches to characterize it. I wrote several sections of the supplement and contributed to the writing of the main text of the paper. Analysis of the crystal structures of ring-like complexes was performed by E. Deeds.

Optimizing ring assembly reveals the strength of weak interactions

Eric J. Deeds^{a,1,2}, John A. Bachman^{b,1}, and Walter Fontana^{b,2}

^aCenter for Bioinformatics and Department of Molecular Biosciences, University of Kansas, Lawrence, KS 66047; and ^bDepartment of Systems Biology, Harvard Medical School, Boston, MA 02115

Edited by José N. Onuchic, University of California at San Diego, La Jolla, CA, and approved December 27, 2011 (received for review August 13, 2011)

Most cellular processes rely on large multiprotein complexes that must assemble into a well-defined quaternary structure in order to function. A number of prominent examples, including the 20S core particle of the proteasome and the AAA+ family of ATPases, contain ring-like structures. Developing an understanding of the complex assembly pathways employed by ring-like structures requires a characterization of the problems these pathways have had to overcome as they evolved. In this work, we use computational models to uncover one such problem: a deadlocked plateau in the assembly dynamics. When the molecular interactions between subunits are too strong, this plateau leads to significant delays in assembly and a reduction in steady-state yield. Conversely, if the interactions are too weak, assembly delays are caused by the instability of crucial intermediates. Intermediate affinities thus maximize the efficiency of assembly for homomeric ring-like structures. In the case of heteromeric rings, we find that rings including at least one weak interaction can assemble efficiently and robustly. Estimation of affinities from solved structures of ring-like complexes indicates that heteromeric rings tend to contain a weak interaction, confirming our prediction. In addition to providing an evolutionary rationale for structural features of rings, our work forms the basis for understanding the complex assembly pathways of stacked rings like the proteasome and suggests principles that would aid in the design of synthetic ring-like structures that self-assemble efficiently.

computational modeling | kinetic deadlock | ring complexes | self-assembly kinetics | glassy dynamics

The vast majority of cellular processes, from signal transduction to the synthesis and degradation of polypeptide chains, rely on the action of large macromolecular complexes (1). In order to carry out their functions, these complexes must adopt a well-defined quaternary structure (1–5). The efficient and effective assembly of these structures from a set of monomeric subunits is thus critically important to living systems. Although experimental work has revealed many details of complex assembly pathways (4–7), conceptual issues remain that are best understood through the analysis of models.

One such issue concerns the evolutionary pressures that have shaped assembly pathways. A similar question has arisen before in the theoretical study of protein folding (8, 9). In that case, it was helpful to consider a “null model,” often called the “Levinthal paradox” (8), which immediately suggested a kinetic problem that protein sequences must overcome in order to fold quickly. Seen from that perspective, evolution has sculpted free energy landscapes that prevent the folding process from degenerating into a random search of conformational space (8, 9) or from producing overly stable intermediates (10–13). In the case of macromolecular assembly, the question revolves not around the free energy landscape that characterizes sequences that fold efficiently, but rather the evolution of the chemical potential landscape of a complex molecular interaction network that supports efficient assembly. If such “assembly landscapes” have been shaped by evolution, what problems have they evolved to overcome or avoid?

In this work, we begin to approach this question for a subset of macromolecular structures; namely, those consisting of rings.

Rings represent a common “motif” in large macromolecular complexes (14), perhaps because of their general thermodynamic stability (15, 16) (see *SI Appendix, Section 1*) and their inherent symmetry. They are thus found in the context of signaling networks (e.g., the apoptosome; refs. 17 and 18), chaperones (e.g., GroEL; ref. 19), protein degradation [e.g., the proteasome (refs. 5 and 7), and ClpP (ref. 20) in bacteria], pore-forming endotoxins [e.g., the protective antigen (PA) of *Bacillus anthracis*; ref. 21], and many other biological processes. Previous studies employing assembly models of ring-like structures have focused on a few specific examples, such as ClpA (22) (an AAA+ family member) and the apoptosome (23). In this work, we focus on a simple but general model of ring assembly, a null model that allows us to illustrate a tension that arises between energetically local interactions and global topological constraints. The barriers induced by this tension can have a strong impact on assembly efficiency, and by understanding how such barriers can be overcome, we provide a basic insight into the evolutionary pressures that have shaped the assembly of a broad class of macromolecular structures.

Our principal finding is the existence of a “deadlocked plateau” in the assembly dynamics of rings and a simple strategy for avoiding it. Depending on the strengths of the molecular interactions between the subunits of the ring, this plateau can have a significant effect on the assembly efficiency of the structure. This is true both for cases in which assembly occurs from an initial condition in which all subunits are monomers, or when considering a steady-state scenario with constant synthesis of monomers and degradation/dilution of complexes. Assembly deadlocks are thus likely to exhibit significant evolutionary pressures on the interaction strengths in the ring. We have also found that, for heteromeric rings where the affinities between neighboring subunits can vary independently, inclusion of one or more “weak” interactions in the structure improves assembly efficiency dramatically. This computational observation leads us to predict that heteromeric ring-like structures will generally contain one interaction that is significantly weaker than the others. We tested this prediction by analyzing the solved structures of all heteromeric three-membered rings, and we found that the vast majority of them do in fact contain at least one weak interaction. Our work thus provides an evolutionary rationale for the structural features of ring-like complexes, in addition to suggesting simple principles that could prove useful in the design of self-assembling nanostructures (24).

Author contributions: E.J.D., J.A.B., and W.F. designed research; E.J.D. and J.A.B. performed research; E.J.D., J.A.B., and W.F. analyzed data; and E.J.D., J.A.B., and W.F. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹E.J.D. and J.A.B. contributed equally to this work.

²To whom correspondence may be addressed. E-mail: deeds@ku.edu or walter@hms.harvard.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1113095109/-DCSupplemental.

Results

Constructing a Model of Ring Assembly. The ring-like protein complexes we model in this work exhibit fairly rigid interaction geometries (such as the structure in Fig. 1A) and a well-defined number n of subunits. Assembly occurs due to binding reactions between intermediates ranging in size from 1 (monomers) to $n - 1$. We account for the geometry of rings by distinguishing three cases for the association of two intermediates with lengths k and l : (i) If $k + l < n$, the binding reaction produces another intermediate species (Fig. 1B); these reactions occur with a uniform association rate α and a dissociation rate β that depends on the strength of the noncovalent bond being formed. (ii) If $k + l = n$, the interaction results in the formation of a ring structure; these reactions involve the (essentially) simultaneous formation of two interfaces (16) (Fig. 1C). Because rings are inherently very stable (15, 16) (see also *SI Appendix, Section 1*), the reverse rate for these reactions (γ) is generally very small. Ring

formation is thus essentially irreversible for most of the parameters considered in this work, and the equilibrium yield of the ring approaches 100%. (iii) If $k + l > n$, the interaction would yield a protein complex with more than n components, resulting in a steric clash that prevents the interaction from occurring (Fig. 1D).

The rings we consider also have a “sidedness”—that is, the monomers in the ring are not themselves internally symmetric; this characterizes most ring-like structures observed in nature (14). As a consequence, the subunits in our model have a distinct left and right side, and interactions can only occur between an interface on the right of one subunit and an interface on the left of another.

Our analysis largely focuses on the simplest case of ring assembly, where the parameters α , β , and γ in Fig. 1 depend only on the identity of the interface(s) involved. In this case, it is straightforward to derive from the reaction classes described above a system of ordinary differential equations (ODEs) describing the assembly dynamics of a ring of length n . The process for deriving these ODEs for both *homomeric* rings (where the individual subunits are indistinguishable) and *heteromeric* rings (where each subunit is distinct from every other subunit) is summarized in *Methods* and is described in detail in the *SI Appendix, Section 2*. The ODEs we obtain are integrated numerically using MATLAB (25) for a given set of parameters (i.e., monomer concentrations, association and dissociation rate constants; see *SI Appendix, Section 2.6*).




Fig. 1. Schematic of ring assembly. (A) A three-membered ring (X-ray structure from Protein Data Bank ID 2JB8) on the left is represented on the right as a graph involving three proteins, each with two binding interfaces (the small circles on the periphery of the nodes). (B) A pair of monomers bind to form a dimer, which represents the case where $k + l < n$. The forward (backward) rate constant of the interaction is denoted α (β). (C) A monomer binds to a dimer yielding the full three-membered ring, which represents the case where $k + l = n$. The forward rate of this reaction is taken to be α as in B, but the backward rate constant (γ) of the reaction is different. Because two interfaces are formed on the right-hand side of the reaction, $\gamma \ll \beta$ (see *SI Appendix, Section 1*). (D) Two dimers attempting to bind—i.e., $k + l > n$. These reactions do not occur because of steric hindrance.

The Assembly Dynamics of Homomeric Rings. Fig. 2A depicts the assembly dynamics obtained from our model of a simple homomeric three-membered ring starting from an initial condition consisting of only monomers. When interaction affinities are very strong, the curves exhibit a characteristic “plateau.” On very short timescales, monomers interact rapidly to form dimers; those dimers can subsequently interact with other monomers to form the full ring. After this initial phase, however, the monomers are depleted from the system but a significant concentration of dimers persists. Because these dimers cannot interact productively either with each other or with the full rings (Fig. 1D), the system is deadlocked until it reaches timescales on which dimers dissociate readily. At that point, monomers released by dissociation can interact with the remaining dimers, resulting in the formation of the full ring. For longer rings, the plateau occurs at lower concentrations of the full ring structure; thus, whereas approximately 65% of three-membered rings are formed in the plateau phase, only approximately 35% of seven-membered rings have formed at that point (see *SI Appendix, Section 4.1.2* and Figs. S6 and S7).

The existence and duration of this deadlocked plateau strongly depends on the parameters of the system. To quantify that dependence, we considered the time T_X it takes a system with an initial condition of 100% monomers (all at equal concentration) to reach a state where $X\%$ of monomers are found in rings. We have plotted T_{99} as a function of the uniform interaction strength along the ring (represented by the dissociation constant for those interactions, $K_d \equiv \beta/\alpha$) for various initial monomer concentrations (Fig. 2B). Each concentration exhibits an affinity that minimizes the time to 99% yield (i.e., T_{99}). Stronger interactions result in considerably longer assembly times: In this “dissociation-challenged” regime, the duration of the deadlocked plateau increases with increasing affinity (see *SI Appendix, Fig. S4*). However, weaker interactions also result in longer assembly times: In this “association-challenged” regime, most dimers do not persist long enough to interact with monomers to create full rings. The value of the optimal affinity is proportional to the concentration of monomers in the system (see *SI Appendix, Fig. S9*).

Although the impact of deadlock on assembly can be quite dramatic, one may ask if this kinetic phenomenon is likely to be important for any given ring. In some cases, such as the apoptosome (23), a ring structure must be populated quickly as a part




Fig. 2. Assembly of a homomeric three-membered ring. (A) In this graph, we consider the percentage of monomers in the various association states (monomer, dimer, and trimer) as a function of time. The affinities are uniformly very strong ($K_d = 10^{-12}$ M). The data are plotted on a logarithmic time-scale because a linear scale obscures the existence of the plateau phase. The on-rate $\alpha = 2.53 \times 10^6 \text{ M}^{-1} \text{ s}^{-1}$ and total subunit concentration $X_T = 400 \text{ nM}$. (B) Variation in assembly time (measured by T_{99} as described in the text) with affinity (K_d) for various initial monomer concentrations X_T . All concentrations exhibit a distinct minimum in T_{99} ; the K_d at which this minimum occurs is proportional to the total monomer concentration (see *SI Appendix, Fig. S9*). α as in A. (C) Steady-state yield (defined as the fraction of monomers in the full ring) as a function of affinity when subunit synthesis and degradation are taken into account according to model A (see *SI Appendix, Sections 2.4 and 4.2*). The synthesis and degradation parameters were chosen to yield the average concentration and half-life of proteins in *Saccharomyces cerevisiae* (28, 29), approximately 480 nM and 42 min, respectively. The solid curve represents an analytical solution of the steady-state yield and the circles represent steady-state results from the numerical integration of model A (see *SI Appendix, Sections 2.4.1 and 3.2.1*). The parameters in this case are α as in A, monomer synthesis rate $Q = 1.31 \times 10^{-10} \text{ M s}^{-1}$, degradation rate $\delta = 2.75 \times 10^{-4} \text{ s}^{-1}$, and $X_T = Q/\delta = 477 \text{ nM}$.

of the propagation of a signal in a signaling network. In that scenario, a plateau could be detrimental because a large fraction of the monomers present in the system may not incorporate into the active molecule on the timescale of the response to signal (23). The interfaces in such rings may thus be under considerable evolutionary pressure to minimize the assembly time of the molecule. However, not all signaling molecules may be sensitive to short-timescale yield. If signaling is functional with the fraction of assembled structures at a level “below” the plateau, then there may be little evolutionary pressure on the affinities in the ring (see *SI Appendix, Sections 4.1.1 and 4.1.3*).

Not all rings in cells may need to assemble quickly; rings are often found as constitutively active and stable assemblies (such as the proteasome or GroEL) that are typically being lost from the cell by active protein degradation and/or dilution arising from cell growth and division. In this case, monomers must be constantly synthesized and assembled into the active structure in order to replace those that are lost. To explore the effect of the phenomena described in Fig. 2B on assembly when accounting for synthesis and degradation, we considered two models. In one case, every monomer in the system has the same probability of being degraded, regardless of the molecular context in which that subunit is found—we term this “model A.” This case represents a likely scenario for active degradation by certain proteases (26, 27). In the second case (“model B”), all complexes have the same probability of being degraded, which corresponds to a situation in which all complexes are being diluted due to rapid cell growth as well as the activity of some proteases (26). In both models, monomers are synthesized at a constant rate; a full description of these models can be found in the *SI Appendix, Section 2.4*.

In this situation, steady-state assembly yield represents essentially the “return on investment” in the energy required for monomer synthesis because monomers that do not incorporate into the active structure are essentially wasted. In Fig. 2C, we plot the steady-state yield of the full complex vs. affinity for a homomeric three-membered ring under model A. The synthesis and degradation parameters in this case were chosen to represent the average concentration and half-life of proteins in *Saccharomyces cerevisiae* (28, 29) (approximately 480 nM and 42 min, respectively). Intermediate affinities maximize yield just as they minimize assembly time, although the magnitude of the effect depends on the parameters. In particular, if degradation rates become very low, the system approaches equilibrium and the greater thermodynamic stability observed for stronger interactions leads to higher yields for those structures (30). The results for model B are similar to those for model A, but with a smaller relative increase in yield (see *SI Appendix, Section 4.2*).

Heteromeric Ring Assembly and the Benefit of Weak Interactions. In heteromeric rings, every single subunit represents a distinct protein. In our models, all of the interactions between proteins along the ring are considered to be specific; that is, a subunit will only bind with its two neighboring proteins and not with any of the other subunits in the ring. When all of the subunit concentrations along the ring are equal, and all of the affinities between subunits are equivalent, one can show that the assembly dynamics of the heteromeric case is actually equivalent to the dynamics of homomeric rings described above (see *SI Appendix, Section 2.3*).

A major difference between homomeric and heteromeric rings, however, is that all of the interaction strengths along a heteromeric ring can be varied independently. We thus examined how changing the relative affinity along the ring influences assembly efficiency by considering a set of seven different affinities ($K_d = 10^{-12}, 10^{-11}, \dots, 10^{-6} \text{ M}$) and constructing all of the unique configurations for a heteromeric ring of length n , where each of the affinities is chosen independently from that set (see *SI Appendix, Section 4.3.1*).

erably smaller than 0.5 (see *SI Appendix, Section 5.4*). Other estimates of affinity (total buried surface area or Protein Interfaces, Surfaces, and Assemblies free energy; ref. 34) yielded similar results (see *SI Appendix, Section 5.3*). Interestingly, our buried surface area results would predict that the “strong” interactions in rings have an average K_d of approximately 10^{-12} M, whereas weak interactions have an average K_d of approximately 10^{-6} M (see *SI Appendix, Section 5.6*), which are precisely the values we employed in Figs. 2*A* and 3*A* and *B*.

Because any heteromeric structure with three interactions will exhibit a weakest and a strongest, we performed two controls to evaluate the significance of the distribution we observed. In the first case, we examined the solved structures of chains. Four-membered chains have the same number of interactions as three-membered rings, but in contrast to rings their assembly efficiency is maximized when all of the interactions are uniformly strong (see *SI Appendix, Sections 3.2.3* and *4.3.3*). We constructed a dataset of 33 structures of heteromeric four-membered chains, whose bimodal W/S distribution is also shown in Fig. 4*A*. Chains have an average ratio (0.63) that is significantly higher than that for rings ($p = 10^{-5}$ based on a random permutation test; see *SI Appendix, Section 5.4*). This difference in average is mostly due to the fact that the weak interaction in chains is, on average, much stronger than the weak interaction in rings ($p = 6 \cdot 10^{-5}$), as we would expect from our findings on assembly efficiency *SI Appendix, Fig. S25*. The majority of chains (22 of the 33) are found in the right peak of the distribution, and in those cases the differences are even more pronounced.

As an additional control, we considered a case in which all of the interactions in the structure were drawn from the same underlying Gaussian distribution (see *SI Appendix, Section 5.5*). Fig. 4*B* shows that both rings and chains exhibit average ratios outside the 95% confidence intervals for this model, indicating that it is unlikely to describe either case. Although we cannot rule out a situation in which affinities are drawn from some other underlying distribution, Fig. 4*B* suggests that the parameters of the distribution could well be under selective pressure to produce rings that meet the affinity requirements for efficient assembly.

Discussion

A number of physical and biological systems, such as glasses and proteins, consist of many concurrently and locally interacting parts. It has long been appreciated that the functional behavior and evolutionary dynamics of these systems are governed by free energy landscapes with many local optima arising from conflicting interactions that are impossible to satisfy simultaneously (e.g., “frustration”; refs. 9 and 30). As was shown for the folding of proteins in the β -trefoil family (10–13), such situations generate a trade-off in which the desirable stability of native contacts (i.e., interactions present in the final configuration) may conflict with the need to undo them should they be generated in the “wrong” temporal order, preventing further native contacts from forming. Prematurely formed native contacts that are too strong have the potential to slow down the required backtracking and significantly delay the overall folding process. Contacts that are too weak, however, destabilize the entire folding process.

In this work, we expand this idea into the realm of assembly, specifically the assembly of rings, where concurrent exploration of all possible assembly pathways leads to an analogous phenomenon, but in the context of a (partially) bimolecular reaction network. Glassy dynamics arises when earlier reactions use up components needed in subsequent reactions (35), thus slowing down the overall kinetics of the final product. Excessive affinity between subunits causes their rapid sequestration into stable intermediates, choking subsequent bimolecular reactions in which these subunits are needed and causing them to be dominated by the dissociation of stable intermediates (corresponding to the “backtracking” in the β -trefoil case). The inclusion of a single

weak interaction in a heteromeric, three-membered ring optimally solves this conundrum by destabilizing only a single intermediate, whose rapid dissociation regenerates monomers ready to react with the other, stable dimers to form the full ring. These results suggest that the chemical potential landscape governing assembly kinetics must evolve features that avoid reaction deadlock, much as free energy landscapes in protein folding must evolve to destabilize certain intermediates in topologically frustrated folds (10–13). Our data analysis of available structures indicates that the “single weak interaction” strategy is likely employed by the majority of evolved heteromeric three-membered rings (Fig. 4*A*). This strategy might serve as a useful guide in the design of synthetic ring-like structures that quickly assemble with high yield (24).

Because assembly arises from a network of bimolecular association and unimolecular dissociation reactions, assembly systems can exhibit features that are not readily observed in the unimolecular isomerization process of protein folding (8, 9). For instance, overexpressing just one subunit of a three-member heteromeric ring severely exacerbates deadlock (see *SI Appendix, Section 4.5*), reinforcing the fact that the operant concern in assembly is a landscape of chemical potential. In addition, assembly systems may employ unique strategies such as subcellular localization of subunits or extensive allosteric interactions among subunits (36) to overcome deadlock. Although our preliminary findings indicate that allostery offers little benefit over the single weak interaction strategy for single rings (see *SI Appendix, Section 4.4*), such approaches may be employed extensively in more complex structures like the proteasome or ribosome (4–7). Our work indicates that the problems of intramolecular folding and intermolecular assembly may share a level of abstraction that enables lessons from landscape theory (9–13), developed in the context of protein folding, to assist in rationalizing the complex assembly mechanisms observed for macromolecular machines.

Methods

Mathematical Model. The mathematical framework we use for modeling the dynamics of ring assembly is explained in detail in the *SI Appendix, Section 2*. We provide a brief description of our approach here. For any homomeric ring of length n , there are n different molecular species that could be generated, ranging from monomers (size 1) to the full ring (size n). The concentration of any species of size j is denoted X_j . For any species of size $j < n$, there are six distinct physical processes that will influence its concentration: (i) an increase in X_j resulting from the dissociation of any larger intermediate that contains it as a subcomplex; (ii) an increase in X_j resulting from a binding interaction between two smaller intermediates; (iii) a decrease in X_j resulting from an interaction with some other intermediate to form a larger complex, but not the full ring; (iv) a decrease in X_j when it dissociates to form smaller intermediates; (v) a decrease in X_j resulting from an interaction with its complementary intermediate to form the full ring; and (vi) an increase in X_j resulting from the dissociation of the full ring.

For the full ring, there are only two processes that affect its concentration: (i) an increase in X_n resulting from a binding reaction between two intermediates, and (ii) a decrease in X_n due to the dissociation of the full ring.

From the processes listed above we can derive a system of ODEs describing the time evolution of the concentration of any intermediate X_j and the full ring X_n (see *SI Appendix, Section 2.1*). Heteromeric rings are modeled in much the same way; the main difference is that there are n distinct molecular species for each size class j (depending on the identities of the subunits in the complex), but only one molecular species for the full ring. Given that only “neighboring” heteromeric intermediates can interact with one another, it is straightforward to derive the ODEs for the heteromeric case (see *SI Appendix, Section 2.2*). We add synthesis and degradation to the model (based either on model A or model B) by including a constant synthesis term (denoted by the variable Q) to the kinetic equation for monomers and the appropriate first-order degradation terms (with a constant degradation rate δ ; see *SI Appendix, Section 2.4*). Our model for homomeric chains is described in the *SI Appendix, Section 2.5*.

All systems of ODEs were numerically integrated using the “ode15s” function in MATLAB (25).

Structural Data. As mentioned in the text, we used the database 3D Complex as a basis for obtaining both the heteromeric three-membered ring and four-membered chain structures (14). In both cases, we chose the “QS-90” level of the 3D Complex hierarchy in order to avoid counting very closely related structures (which are often simply mutants of a single protein) in the dataset. Of the 82 rings in this set, many represent situations quite distinct from that considered in our model. For instance, antibody–antigen complexes often form three-membered rings (involving the heavy and light chains of the antibody, which bind each other and the antigen), but such structures have not evolved to assemble with all three chains present. Rather, the antibody chains are assembled first in cells, and only when secreted (or expressed on a cell surface) do they interact with the antigen. Similarly, a number of “three-membered” rings in 3D Complex involve proteases bound to a protein inhibitor. In those cases, the two chains of the protease are actually synthesized as a long polypeptide chain that is cleaved during maturation of the zymogen. The interaction between these chains thus does not arise as a result of a bimolecular reaction, but rather a unimolecular folding reaction, and as such the assembly of these structures is not considered in our model. In total, 53 of the 82 three-membered rings were deemed to not conform to the assumptions of our model, leaving 29 structures for the analysis in Fig. 4. Similarly, of the 104 four-membered heteromeric chains we obtained from 3D Complex, 60 were disregarded for reasons similar to the ones cited for rings, and 11

were actually found to be rings on further analysis. A full description of the datasets and their construction can be found in *SI Appendix, Sections 5.1 and 5.2*. A detailed list of all structures included in the datasets, or excluded for one of the reasons cited, is also provided in the *SI Table of Structures*.

Statistical Methods. To test if the affinity distributions we observed exhibited significantly different averages, we performed a simple permutation test using the “twot.permutation” function provided by the Data Analysis and Graphics (DAAG) package in R (37) with 10^5 replicates. The p value reported represents the fraction of these permuted datasets with a difference of means greater than the difference we observed. The Gaussian control in Fig. 4B was obtained by sampling three affinities from an underlying Gaussian distribution with an average and standard deviation similar to that observed for both our dataset of three-membered rings and our dataset of four-membered chains. A more detailed description of the affinity distributions can be found in the *SI Appendix, Sections 5.4 and 5.5*.

ACKNOWLEDGMENTS. We would like to thank Tom Kolokotronis for his suggestions and his comments on the manuscript. E.J.D. was supported by a postdoctoral fellowship from the National Institutes of Health.

- Alberts B (1998) The cell as a collection of protein machines: Preparing the next generation of molecular biologists. *Cell* 92:291–294.
- Korostelev A, Noller HF (2007) The ribosome in focus: New structures bring new insights. *Trends Biochem Sci* 32:434–441.
- Bashan A, Yonath A (2008) Correlating ribosome function with high-resolution structures. *Trends Microbiol* 16:326–335.
- Staley JP, Woolford JL (2009) Assembly of ribosomes and spliceosomes: Complex ribonucleoprotein machines. *Curr Opin Cell Biol* 21:109–118.
- Murata S, Yashiroda H, Tanaka K (2009) Molecular mechanisms of proteasome assembly. *Nat Rev Mol Cell Biol* 10:104–115.
- Kressler D, Hurt E, Bassler J (2010) Driving ribosome assembly. *Biochim Biophys Acta* 1803:673–683.
- Marques AJ, Palanimurugan R, Matias AC, Ramos PC, Dohmen RJ (2009) Catalytic mechanism and assembly of the proteasome. *Chem Rev* 109:1509–1536.
- Shakhnovich E (2006) Protein folding thermodynamics and dynamics: Where physics, chemistry, and biology meet. *Chem Rev* 106:1559–1588.
- Onuchic JN, Wolynes PG (2004) Theory of protein folding. *Curr Opin Struct Biol* 14:70–75.
- Capraro DT, Roy M, Onuchic JN, Jennings PA (2008) Backtracking on the folding landscape of the beta-trefoil protein interleukin-1beta? *Proc Natl Acad Sci USA* 105:14844–14848.
- Chavez LL, Gosavi S, Jennings PA, Onuchic JN (2006) Multiple routes lead to the native state in the energy landscape of the beta-trefoil family. *Proc Natl Acad Sci USA* 103:10254–10258.
- Gosavi S, Chavez LL, Jennings PA, Onuchic JN (2006) Topological frustration and the folding of interleukin-1 beta. *J Mol Biol* 357:986–996.
- Gosavi S, Whitford PC, Jennings PA, Onuchic JN (2008) Extracting function from a beta-trefoil folding motif. *Proc Natl Acad Sci USA* 105:10384–10389.
- Levy ED, Pereira-Leal JB, Chothia C, Teichmann SA (2006) 3D complex: A structural classification of protein complexes. *PLoS Comput Biol* 2:e155.
- Saiz L, Vilar JMG (2006) Stochastic dynamics of macromolecular-assembly networks. *Mol Syst Biol* 2:006.0024.
- Bray D, Lay S (1997) Computer-based analysis of the binding steps in protein complex formation. *Proc Natl Acad Sci USA* 94:13493–13498.
- Qi S, et al. (2010) Crystal structure of the *Caenorhabditis elegans* apoptosome reveals an octameric assembly of CED-4. *Cell* 141:446–457.
- Shi Y (2006) Mechanical aspects of apoptosome assembly. *Curr Opin Cell Biol* 18:677–684.
- Braig K, et al. (1994) The crystal structure of the bacterial chaperonin GroEL at 2.8 Å. *Nature* 371:578–586.
- Wang J, Hartling JA, Flanagan JM (1997) The structure of ClpP at 2.3 Å resolution suggests a model for ATP-dependent proteolysis. *Cell* 91:447–456.
- Petosa C, Collier RJ, Klimpel KR, Leppla SH, Liddington RC (1997) Crystal structure of the Anthrax toxin protective antigen. *Nature* 385:833–838.
- Kress W, Mutschler H, Weber-Ban E (2007) Assembly pathway of an AAA+ protein: Tracking ClpA and ClpAP complex formation in real time. *Biochemistry* 46:6183–6193.
- Nakabayashi J, Sasaki A (2006) A mathematical model for apoptosome assembly: The optimal cytochrome c/Apaf-1 ratio. *J Theor Biol* 242:280–287.
- Yin P, Choi HMT, Calvert CR, Pierce NA (2008) Programming biomolecular self-assembly pathways. *Nature* 451(7176):318–322.
- MATLAB (2010) MATLAB Ver 7.11. (MathWorks, Natick, MA).
- Sharma S, Hoskins JR, Wickner S (2005) Binding and degradation of heterodimeric substrates by ClpAP and ClpXP. *J Biol Chem* 280:5449–5455.
- Moore SD, Baker TA, Sauer RT (2008) Forced extraction of targeted components from complex macromolecular assemblies. *Proc Natl Acad Sci USA* 105:11685–11690.
- Belle A, Tanay A, Bitincka L, Shamir R, O’Shea EK (2006) Quantification of protein half-lives in the budding yeast proteome. *Proc Natl Acad Sci USA* 103:13004–13009.
- Ghaemmaghami S, et al. (2003) Global analysis of protein expression in yeast. *Nature* 425:737–741.
- Wolynes PG (1996) Symmetry and the energy landscapes of biomolecules. *Proc Natl Acad Sci USA* 93:14249–14255.
- Horton N, Lewis M (1992) Calculation of the free energy of association for protein complexes. *Protein Sci* 1:169–181.
- Bougouffa S, Warwicker J (2008) Volume-based solvation models out-perform area-based models in combined studies of wild-type and mutated protein-protein interfaces. *BMC Bioinformatics* 9:448.
- Cavallo L, Kleijung J, Fraternali F (2003) POPS: A fast algorithm for solvent accessible surface areas at atomic and residue level. *Nucleic Acids Res* 31:3364–3366.
- Krissinel E, Henrick K (2007) Inference of macromolecular assemblies from crystalline state. *J Mol Biol* 372:774–797.
- Awazu A, Kaneko K (2009) Ubiquitous “glassy” relaxation in catalytic reaction networks. *Phys Rev E* 80:041931.
- Williamson JR (2008) Cooperativity in macromolecular assembly. *Nat Chem Biol* 4:458–465.
- R Development Core Team (2010) *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, Vienna).

Optimizing ring assembly reveals the strength of weak interactions

Supporting Information Appendix

Eric J. Deeds^{1*}, John A. Bachman^{2*} and Walter Fontana²

¹Center for Bioinformatics and Department of Molecular Biosciences, The University of Kansas, 2030 Becker Dr., Lawrence, KS 66047

²Department of Systems Biology, Harvard Medical School, 200 Longwood Avenue, Boston MA 02115, USA

Email: Eric Deeds - deeds@ku.edu; John Bachman - bachman@fas.harvard.edu; Walter Fontana - walter@hms.harvard.edu;

*These authors contributed equally

Contents

List of Figures	3
Section summaries	4
1 Thermodynamics of rings	6
2 Mathematical model of ring assembly	9
2.1 Assembly of homomeric ring complexes	9
2.1.1 Notation	9
2.1.2 Structure, symmetries, and rate constants	9
2.1.3 Ordinary differential equations	12
2.1.4 Example: homomeric three-ring	14
2.2 Assembly of heteromeric rings	14
2.2.1 Notation	14
2.2.2 Ordinary differential equations	15
2.2.3 Example: heteromeric three-ring	17
2.3 Homomeric ring assembly arises as a special case of heteromeric ring assembly . . .	17
2.4 Modeling synthesis and degradation of rings	22
2.4.1 Homomeric model A: subunit deletion	23
2.4.2 Heteromeric model A: subunit deletion	25

2.4.3	Homomeric model B: whole-complex degradation	27
2.4.4	Heteromeric model B: whole-complex degradation	28
2.5	Assembly of Heteromeric Chains	30
2.6	A note on numerical methods	32
3	Equilibrium and steady-state solutions for the homomeric 3-membered ring	32
3.1	Equilibrium results without synthesis or degradation	32
3.2	Steady-state results including synthesis and degradation	33
3.2.1	Model A: subunit deletion	33
3.2.2	Model B: whole-complex degradation	35
3.2.3	Model A degradation for a homomeric chain	36
4	Additional results	38
4.1	Ring assembly dynamics and deadlock	38
4.1.1	Effects of affinity and concentration on ring assembly	38
4.1.2	Scaling of plateau height with ring length	39
4.1.3	Assembly time as a function of affinity, ring length, and concentration	41
4.2	Synthesis and degradation	42
4.2.1	Ring assembly dynamics with synthesis and degradation	42
4.2.2	Effect of K_D and synthesis rate on assembly yield	44
4.3	Effect of affinity configurations on assembly	47
4.3.1	Enumerating distinct affinity configurations for heteromeric rings	47
4.3.2	Numerical simulation results for the heteromeric 3-, 4-, 5-, and 6-rings	49
4.3.3	Optimizing the assembly of chains	51
4.4	Hierarchical assembly pathways	52
4.4.1	Mathematical model of stepwise assembly	53
4.4.2	Comparing sequential assembly with weak interactions	53
4.5	Non-uniform concentrations	56
5	Analysis of structural data	57
5.1	Structures for heteromeric three-membered rings	57
5.2	Structures for heteromeric four-membered chains	58
5.3	Calculating changes in non-polar surface area	60

5.4	Comparing rings and chains	60
5.5	Gaussian control	64
5.6	Affinities for the interactions in the crystal structures of rings	65
References		66

List of Figures

1	Schematic free energy landscape	8
2	Structured reactions require explicit combinatorial factors	10
3	The equilibrium fraction of monomers and trimers for a 3-ring	34
4	Effect of affinity on ring assembly dynamics	38
5	Effect of initial monomer concentration on ring assembly dynamics	39
6	Effect of ring length on ring assembly dynamics	40
7	Effect of ring length on plateau height	40
8	Assembly time as a function of affinity, concentration, and ring length	42
9	Optimal affinity as a function of monomer concentration	43
10	Ring assembly in the presence of synthesis and degradation	43
11	Assembly dynamics vs.affinity in the presence of synthesis and degradation	44
12	Steady state yield vs.affinity with synthesis and degradation	45
13	Steady state yield vs.affinity and total protein concentration (model A)	46
14	Assembly yield at optimal affinity vs.strong interactions for different X_T	46
15	Ring symmetries	48
16	Ranking the assembly efficiency and yield of affinity configurations for the 3-ring	49
17	Ranking the assembly efficiency and yield of affinity configurations for the 4-ring	50
18	Ranking the assembly efficiency and yield of affinity configurations for the 5-ring	50
19	Ranking the assembly efficiency and yield of affinity configurations for the 6-ring	51
20	Ranking the assembly efficiency and yield of affinity configurations for the 4-chain	52
21	3-rings with uniform interactions, one weak interaction, and sequential assembly	54
22	Schematic diagram of ring assembly	55
23	Assembly dynamics of a 3-ring with non-uniform subunit concentrations	56
24	Strong interactions in rings and chains	61
25	Weak interactions in rings and chains	62

26	The weak-to-strong ratio for rings and chains	62
27	Comparing the ratio distributions of rings and chains with normal distributions . . .	63
28	Change in solvent-accessible non-polar surface area and free energy of binding	65

Section summaries

1—Thermodynamics of rings

In this section we provide a brief explanation for the fact that rings are typically orders of magnitude more stable than acyclic intermediates.

2—Mathematical model of ring assembly

In sections 2.1 and 2.2 we provide a detailed derivation of the kinetic equations for the assembly mechanism that defines our model of homo- and heteromeric ring formation. Some care is required to properly account for combinatorial factors when formulating a model in terms of monomers with explicit interfaces with elementary interactions and associated rate constants.

In section 2.3 we prove that the equations for homomeric ring assembly are a rescaled version of those for heteromeric ring assembly when all affinities are equal and all intermediates in a size class have equal initial concentration.

In section 2.4 we extend our model of ring assembly to include synthesis and degradation. We consider two variants, A and B, that differ in the degradation mechanism. In model A, each *monomer* is removed at a rate proportional to its concentration, implying that the degradation machinery can pry individual monomers out of a complex (including the full ring). In model B, each *complex* is removed at a rate proportional to its concentration.

Since they lack a global geometric constraint, the assembly of chains provides an insightful contrast to the assembly of rings, not the least because a weak interaction makes the ring case more similar to the chain case, while retaining the stability intrinsic to rings. In section 2.5 we define the kinetic equations for assembly and synthesis/degradation of type A and B for heteromeric chains of four monomers. (We discuss chain optimization in section 4.3.3.)

3—Equilibrium and steady-state solutions for the homomeric 3-membered ring

In section 3.1 we provide an analytical solution for the concentration of the three-membered ring at equilibrium based on the thermodynamic considerations of section 1. Even for comparatively weak interactions, small amounts of monomers suffice to push the equilibrium almost entirely to the side of rings, explaining why the yields we observe are very close to 100% (see also Figure 2A of the main text) in the parameter regimes we consider.

In section 3.2.1 we derive the equation for the steady-state yield of the three-member ring under type-A degradation shown in Figure 2C of the main text. In section 3.2.2 we derive the equation for the steady-state curve of 3-ring assembly under type-B degradation. We compare type A and type B models in section 4.2.2 (see Figure 12).

In section 3.2.3 we show that, in contrast to the ring case, strengthening the interactions of chains under a type A degradation scenario always increases yield (see sections 3.2.1 and 3.2.2; see also

Figure 2C of the main text).

4—Additional results

In section 4.1 we examine the formation of the plateau phase for 3-rings as a function of affinity and concentration. The plateau begins to appear at a critical interaction strength (depending on total monomer concentration) and increases in duration with increasing affinity, while its height remains unaffected. The plateau height, however, decreases with increasing ring size. We also examine the chevron plots of Figure 2B in the main text as a function of ring size and target yields.

In section 4.2 we compare in more detail the assembly kinetics of the degradation models A and B described in section 2.4. We examine the effect of interaction strength and synthesis rate on yield for models A and B using the equations derived in sections 3.2.1 and 3.2.2. Interactions that are weaker or stronger than an optimal value reduce steady-state assembly yield.

In section 4.3 we assess the effect of affinity configurations on ring assembly with respect to efficiency (time to 99% equilibrium yield or steady-state yield for the two versions with degradation). We first detail how to properly enumerate distinct affinity configurations and then sample configurations for heteromeric 4-, 5-, and 6-rings, extending the results for 3-rings reported in the main text. The main observation is that rings with two weak interactions perform well, presumably because they can assemble as two “quasi-chains” and then snap together. We compare these results with the case of heteromeric chain assembly, showing that universally stronger interactions are always better.

While heterogeneity of affinities is one solution to assembly bottlenecks, a strict sequential (“hierarchical”) assembly—based, for example, on allosteric interactions—might be another, as touched upon in the main text. In section 4.4 we compare three assembly scenarios (uniform interactions, one weak interaction, hierarchical), suggesting that the weak interaction scenario yields the fastest assembly.

In section 4.5 we consider cases in which the concentration of subunits varies in a three-membered heteromeric ring.

5—Analysis of structural data

In this section we describe the structures in the dataset and the statistical procedures underlying their analysis, summarized as Figure 4 in the main text.

1 Thermodynamics of rings

In order to construct a model of ring assembly, we must first have a general characterization of the chemical reactions that are involved in their formation. For simplicity, we begin by considering the case of a single, homomeric ring of length n . As discussed in the main text, if two intermediates, consisting of k and l subunits respectively, react with one another, there are essentially three classes of reaction:

1. $k + l < n$

The intermediates react and produce another intermediate of length $k + l$ and a single non-covalent interaction is formed.


2. $k + l = n$

The reaction produces a fully formed ring by creating two non-covalent interactions.

3. $k + l > n$

In our models, we imagine that the molecules in question are fairly rigid with well-defined interaction angles. Thus, if $k + l > n$, formation of the product of the reaction would result in two subunits occupying the same space (see Fig. 1D in the main text). We thus ignore reactions of type 3.

Type 1 reactions involve the formation of a single non-covalent interaction in a reversible manner:




where “ X_k ” denotes the chemical species corresponding to an intermediate with k subunits. We denote the forward rate of this reaction (or the “on” rate) as α , and the backward (or “off”) rate as β . In our work, we do not consider small variations in the on rate that might arise for intermediates of differing size, and we take α to be a constant for all the type 1 reactions in the system. We can write the standard free energy of formation of the interaction in type 1 reactions as:

$$\Delta G_b^0(1) = \Delta G_i^0 + \Delta G_p^0, \quad (1)$$

where $\Delta G_b^0(1)$ is the overall standard free energy of formation of the interaction. This free energy is related in the usual way to the dissociation constant of the reaction by $K_D(1) = \beta/\alpha = c_0 e^{\Delta G_b^0(1)/RT}$, where c_0 is some standard reference concentration (here taken to be 1 molar). In equation 1, we have decomposed this binding energy into two components, adopting the notation of Saiz and Vilar [1]. ΔG_i^0 represents the free energy of formation of the interface itself and includes contributions from both the desolvation of the two protein interfaces and the molecular contacts (e.g. electrostatic and Van der Waals interactions) formed upon binding. ΔG_p^0 represents the positional entropy loss entailed when taking two proteins that can freely diffuse around a particular molar volume and confining them to a given binary complex [1].

Type 2 reactions can be written in a very similar manner to type 1 reactions above. We have:



In this case, we assume that the forward rate of reaction is again equal to α and we introduce a third parameter, γ , to represent the dissociation rate for a fully formed ring. Since type 2 reactions constrain the motion of exactly the same number of molecules as type 1 reactions, the change in positional entropy for both reactions is equivalent [1]. A type 2 reaction, however, involves the formation of 2 interfaces, rather than just one (see Fig. 1C in the main text). If we consider a case where all of the interactions along the ring have identical thermodynamic properties, we have:

$$\begin{aligned}\Delta G_b^0(2) &= 2\Delta G_i^0 + \Delta G_p^0 \\ &= 2\Delta G_b^0(1) - \Delta G_p^0\end{aligned}\tag{2}$$

where the second equation relates the energetics of type 2 reactions to the energetics of type 1 reactions. Since the association rate α is taken to be equivalent for both cases here, equation 2 indicates that the reverse rate of a type 2 reaction (γ) will differ significantly from β . Defining $K_\gamma \equiv \gamma/\alpha$ as the dissociation constant for reactions of type 2, we have:

$$\begin{aligned}\gamma &= \alpha \cdot c_0 \cdot e^{(2\Delta G_b^0(1) - \Delta G_p^0)/RT} \\ &= \left(\frac{\alpha}{c_0}\right) K_D(1)^2 \cdot e^{-\Delta G_p^0/RT}.\end{aligned}\tag{3}$$

If all of the individual reactions of type 1 are favorable for a given ring (i.e. $K_D < 1$ M for all reactions of type 1)—and since $\Delta G_p^0 > 0$ by definition—we will have $\gamma \ll \beta$. Thus, for the types of rings discussed in this work, the rate of dissociation for the full ring is generally many orders of magnitude smaller than the rate of dissociation for any of the acyclic intermediates.

We can define similar equations for cases in which the interactions in a ring do not have the same strengths (e.g. the heteromeric cases described in section 2.2 below). In our notation, each interaction in such a context has a unique label (see section 2.2.1); we define the type 1 affinity of any given interaction “ i ” as $K_D(i) \equiv \frac{\beta_i}{\alpha}$, where β_i denotes the off rate for interaction i . Hence in any type 2 reaction in which interactions “ i ” and “ j ” are formed, we have:

$$\gamma_{i,j} = \left(\frac{\alpha}{c_0}\right) K_D(i) \cdot K_D(j) \cdot e^{-\Delta G_p^0/RT}.\tag{4}$$

In our analysis of ring assembly we treat the affinities of the interactions along the ring as our principal variable and examine how changes in affinities influence assembly efficiency. As indicated above, we treat α as a constant—that is, variations in the dissociation constant (either between configurations or across the interactions in a given ring) are taken to represent changes in β and γ (see Figure 1 for an example of this thermodynamic picture for two different dimers forming along a ring). Changing the value of α represents an alternative method of modulating

the affinity. Interestingly, changes in α can be mirrored by changes in concentration (assuming that all reactions still have the same on rate). This is perhaps easiest to see by considering the stochastic version of the deterministic association rate. If we define this stochastic rate as “ a ”, we have that $\alpha = a \cdot N_A \cdot V$ where N_A is Avogadro’s number and V is the volume of the system [2]. In this construction, we note that any change to α can be conceptualized as either a change in the fundamental frequency of collisions (a) while keeping concentration constant, or as a change in the volume of the system (while keeping total particle numbers constant). The latter scenario corresponds to a change in concentration. Thus, changing K_D by varying α is equivalent to changing concentration scales while keeping α constant. The effect of concentration on assembly dynamics can be seen in section 4.1.1, Figure 5.




Figure 1: Schematic free energy landscape for a case in which differences in affinities are entirely represented by differences in off rates. Here we have two different binding reactions: A binds B and C binds D. “A + B” and “C + D” represent the unbound states on the far left of the schematic reaction coordinate; the unbound states in this case have roughly the same free energy. The transition states (represented by “A ... B” and “C ... D”) also have approximately the same free energy; the change in free energy from the unbound state to the transition state is identical in both cases (giving identical values of α). However, the bound states (“AB” and “CD”) exhibit very different free energies, and the difference in free energy change between the transition state and the bound state results in a much higher value of β for the C-D binding reaction compared to the A-B binding reaction.

In this work we take a “ballpark” value for the association rate of proteins as $\alpha \approx 10^6 \text{ M}^{-1} \text{ s}^{-1}$ [3].

Since we take α to be constant, for any given value of K_D , if we know ΔG_p^0 , we can calculate both β and γ for any particular temperature T . Here we focus on $T = 300 \text{ K}$, so $RT \approx 0.6 \text{ kcal mol}^{-1}$. Estimates of ΔG_p^0 vary in the literature [1], but here we set $\Delta G_p^0 = 9 \text{ kcal mol}^{-1}$ [4], which is taken to be a constant for all reactions. Changing the value of ΔG_p^0 will essentially modulate the stability of rings compared to acyclic intermediates, although even considerably smaller values (say, 6 kcal mol^{-1}) still result in $\gamma \ll \beta$ for the parameter values we consider here (see equation 3). We leave investigation of the effects of modulating the positional entropy term to future work.

2 Mathematical model of ring assembly

2.1 Assembly of homomeric ring complexes

2.1.1 Notation

We first consider a homomeric ring structure consisting of n identical subunits, any of which can bind to any other. The subunits are treated as identical but have a “sidedness,” that is, each subunit has two distinct interfaces, a left and a right; the left side of one can bind to the right side of another, while two left sides and two right sides cannot bind each other.

As in section 1, we adopt the notation that X_j denotes the sub-complex containing j subunits, with $1 \leq j \leq n$. In the discussion that follows, we also use X_j to refer to the *concentration* of the complex with j subunits—the meaning of “ X_j ” in any given case will be clear from the context. In the differential equations that follow, X_1 thus denotes the concentration of monomers, X_2 denotes the concentration of dimers, and so on, with X_n denoting the concentration of the full ring structure.

As discussed in section 1, we denote the on rate as α , the off rate as β , and the ring breakage rate (the off rate for type 2 reactions) as γ .

2.1.2 Structure, symmetries, and rate constants

A model of assembly must necessarily take into account essential structural aspects of the molecular species that are generated in the process. This has consequences for how we write our rate equations.

The kinetic description of a reaction splits into a time-dependent monomial of concentrations representing mass-action and a time-independent term, the rate constant, representing a reaction *mechanism*. Models of assembly, like ours, aim at studying the consequences of mechanisms and are therefore defined in terms of thermodynamically motivated rate constants that pertain to elementary interactions (see section 1). This requires that we account explicitly for combinatorial factors intrinsic to a reaction mechanism rather than absorbing them into an overall rate constant.

To illustrate the issue, consider the dimerization of a monomer. The formation of an asymmetric dimer will occur at twice the rate than the formation of a symmetric dimer. Yet, this distinction cannot be expressed in a notation that is too terse in structural detail, such as $A + A \rightarrow A_2$. The situation clarifies instantly when making a minimum of structure explicit. For example, consider identical monomers, each with two distinct binding sites, black and white. The asymmetric case might be represented as $\square\bullet + \square\bullet \rightarrow \square\bullet\square\bullet$ and the symmetric case as $\square\bullet + \square\bullet \rightarrow \square\bullet\bullet\square$. The difference in the respective reaction rates is accounted for by noting that the mechanism producing the asymmetric dimer allows for two “reaction paths”, each involving a distinct black/white combination, while the mechanism producing the symmetric dimer admits only one reaction path (since the “white” sites in the symmetric case cannot bind one another). If we wish to ascribe the same fundamental reaction rate constant to any given “site-site” interaction for both cases, the “apparent” rate constant for the asymmetric reaction will be twice that of the symmetric reaction. By representing molecular structure that is relevant to the problem at hand, we can thus reason about combinatorial factors intrinsic to a mechanism.

Since combinatorial factors apply per reaction event, we must also account for factors that arise from the distinguishability of reactants. Say we have a system of reactions: $A + A \rightarrow A_2$ and $A + A_2 \rightarrow A_3$. In the first case, the reactants are *indistinguishable* from one another, while in the latter case the reactants are distinguishable. If we wish to have the same per-site rate constant for these reactions (as in section 1) then we must consider the number of distinct instances of each reaction. For A binding to A_2 , we will have $N(A) \cdot N(A_2)$ such instances, where $N(A)$ is the number of A molecules in the system; for A binding to itself, we will have $N(A)(N(A) - 1)/2$ distinct reaction instances. In the second case, passing to the limit of continuous concentrations $[A]$, we can neglect the linear term, but must preserve the factor $1/2$. Taking explicitly into account the combinatorial factors arising from *both* mechanism and reactant indistinguishability permits the formulation of rate equations in terms of rate constants that are taken to be the same across reactions. In our example, the asymmetric reaction has rate $v = \alpha \cdot 2 \cdot [A]^2/2 = \alpha \cdot [A]^2$, while the symmetric case has half that rate: $v = \alpha \cdot [A]^2/2$; in both cases $d[A]/dt = -2v$ and $d[A_2]/dt = v$ due to the fact that the reaction “consumes” two molecules of A but produces only one molecule of A_2 .

An analogous situation arises with the assembly of homomeric rings when two intermediates of equal length k react with one another. The asymmetric case occurs when the combined length $2k$ of the product is less than the length n of the ring (Figure 2A). The symmetric case occurs when $k = n/2$ and is due to the simultaneous formation of two interactions upon ring closure (Figure 2B), as argued in section 1.




Figure 2: Structured reactions provide a rationale for explicit combinatorial factors needed for rate equations employing thermodynamic rate constants. The figure illustrates the asymmetric and symmetric dimerization reactions that arise in homomeric ring assembly. The grey boxes represent identical chains of equal length $k > 1$. Asymmetric dimer formation (A) arises when chains of equal length combine to form a chain twice as long but shorter than the full ring. Symmetric dimerization (B) arises only for rings of even length and corresponds to identical chains, each half the ring size, simultaneously forming two interactions and completing the ring, as detailed in section 1. As discussed in the text, the reaction velocity of the asymmetric case is twice the velocity of the symmetric case.

For the sake of less cluttered equations, we will not name molecular species in a manner that fully describes their relevant structure, but use a “shorthand” notation instead as described in section 2.1.1. The drawback of such a notation is that one can forget the implications that structural details can have for determining reaction rates. It is therefore useful to enumerate the various reaction classes and their distinct kinetic contribution to homomeric ring assembly. We essentially have bimolecular association reactions in which intermediates form other intermediates or the full ring, and unimolecular dissociation reactions in which molecules break apart to form smaller intermediates. We denote the kinetic term that a reaction class contributes to species X_k with D_k .

1. **The bimolecular association reaction $X_k + X_l \rightarrow X_j$, with $j < n$.** Since every chain X_k has two “sides,” there are two paths for this reaction to proceed: X_l binds on the right of X_k and X_l binds on the left of X_k . Both paths occur with rate α , yielding a reaction rate $v = 2 \cdot \alpha \cdot X_k \cdot X_l$ with $D_k = D_l = -v$ and $D_j = v$.
2. **The bimolecular association reaction $2 X_k \rightarrow X_j$, with $j < n$.** This represents two copies of a molecule of length k binding to form an intermediate of length $j = 2 \cdot k$. Since one of the X_k ’s can bind on either side of the other (Figure 2A), the rate of reaction is $v = 2 \cdot \alpha \cdot X_k^2/2 = \alpha \cdot X_k^2$ with $D_k = -2v$ and $D_j = v$.
3. **The bimolecular association reaction $X_k + X_l \rightarrow X_n$.** In this case, there is only one reaction path for X_l to bind to X_k : if the right side of X_k binds to the left side of X_l , then the left side of X_k must simultaneously bind the right side of X_l in order to form the full ring (see Figure 1C in the main text). The rate for this reaction is $v = \alpha \cdot X_k \cdot X_l$, with $D_k = D_l = -v$ and $D_n = v$.
4. **The bimolecular association reaction $2 X_k \rightarrow X_n$.** This represents two copies of a molecule of length $k = n/2$ binding to one another to form the full ring (Figure 2B); this reaction only applies to rings of even length. As discussed above, the rate of this reaction is $v = \alpha \cdot X_k^2/2$ with $D_k = -2v$ and $D_n = v$.
5. **The dissociation reaction $X_j \rightarrow X_k + X_l$ with $j < n$.** This represents the dissociation of a chain X_j into two distinct smaller chains X_k and X_l . The dissociation rate of any single interaction is denoted β , but in this case there are always *two* interactions in X_j that could break in order to give X_k and X_l ; in other words, X_k can break off from either the right or left end of X_j . The rate of the reaction is thus $v = 2 \cdot \beta \cdot X_j$ with $D_j = -v$ and $D_k = D_l = v$.
6. **The dissociation reaction $X_j \rightarrow 2 X_k$ with $j < n$.** This represents the dissociation of a chain X_j into two copies of X_k . Of course, j must be even in this case. There is always only one interaction in X_j that can break in order to form 2 X_k ’s (namely, the interaction in the “middle” of X_j) so the rate for this reaction is just $v = \beta \cdot X_j$ with $D_j = -v$ and $D_k = 2v$.
7. **The dissociation reaction $X_n \rightarrow X_k + X_l$.** This represents the dissociation of the full ring into two different intermediates of length k and l . As discussed in section 1, any two interactions in X_n will break with rate γ . Since there are always n ways to choose two interactions in X_n that will form X_k and X_l when both are broken, the rate of this reaction is $v = n \cdot \gamma \cdot X_n$ with $D_n = -v$ and $D_k = D_l = v$.
8. **The dissociation reaction $X_n \rightarrow 2 X_k$.** This represents the dissociation of the full ring into two copies of the intermediate X_k ; this reaction obviously only applies to rings of even length. In this case, there are only $n/2$ distinct ways of choosing two interactions to break in order to form two copies of X_k , so the rate of the reaction is $v = n/2 \cdot \gamma \cdot X_n$ with $D_n = -v$ and $D_k = 2v$.

In the following section, we aggregate over these types of reactions in order to construct a full system of ODEs describing the deterministic time evolution of the concentrations of all molecular species.

2.1.3 Ordinary differential equations

Assembly Intermediates. Assuming mass-action kinetics as described in section 2.1.2 above, the ordinary differential equations describing the change in concentration of any ring-assembly intermediate X_j , $1 \leq j < n$, can be written as follows:

$$\frac{dX_j}{dt} = M_j^{(1)} + M_j^{(2)} + M_j^{(3)} + M_j^{(4)} + M_j^{(5)} + M_j^{(6)} \quad 1 \leq j < n, \quad (5)$$

where the $M_j^{(i)}$ represent groups of terms describing the exchange of mass between intermediates resulting from six specific binding and unbinding processes. The processes are:

1. An increase in X_j resulting from the dissociation of higher-order intermediates containing X_j as a sub-complex (positive terms):

$$M_j^{(1)} = 2 \cdot \beta \sum_{l=j+1}^{n-1} X_l.$$

Note that this term covers both reactions in which only one molecule of X_j is produced (class 5 reactions in section 2.1.2) and reactions where 2 molecules of X_j are produced (class 6 reactions in section 2.1.2).

2. An increase resulting from smaller intermediates binding together to form X_j (positive terms):

$$M_j^{(2)} = \alpha \sum_{l=1}^{j-1} X_l \cdot X_{j-l}.$$

Note that the sum in the above term will count asymmetric pairs of intermediates twice, but will count symmetric pairs only once. To illustrate this, suppose we have $j = 6$ (with of course $n > 6$). The sum in $M_j^{(2)}$ will contain the following terms: $X_1 \cdot X_5$, $X_2 \cdot X_4$, $X_3 \cdot X_3$, $X_4 \cdot X_2$ and $X_5 \cdot X_1$. Every type of interaction occurs twice except for the $j/2$ case. $M_j^{(2)}$ thus correctly represents the difference between class 1 reactions (where the two intermediates are of different length) and class 2 reactions (where the two intermediates have the same length), as described in section 2.1.2.

3. A decrease resulting from the binding of X_j to complementary intermediates to form higher-order complexes, but not the full ring (negative terms):

$$M_j^{(3)} = -2 \cdot \alpha \cdot X_j \sum_{l=1}^{n-j-1} X_l. \quad (6)$$

Note that, unlike $M_j^{(2)}$, symmetric and asymmetric cases are considered equally here. This is due to the fact that both class 1 and class 2 association reactions have the same net effect on the concentration of the reactants (see section 2.1.2).

4. A decrease resulting from the breakage of one of the $j - 1$ interactions within the complex X_j (negative term):

$$M_j^{(4)} = -\beta \cdot X_j \cdot (j - 1).$$

This term covers both class 5 and class 6 dissociation reactions (see section 2.1.2).

5. A decrease resulting from the binding of X_j to its complementary intermediate to form the full ring (negative term):

$$M_j^{(5)} = -\alpha \cdot X_j \cdot X_{n-j}.$$

Note that when X_j interacts with its complement X_{n-j} to form the full ring X_n , *two* new interactions are formed simultaneously (see section 1 and Fig. 1C in the main text). This term covers both class 3 and class 4 association reactions in section 2.1.2, since both have the same net effect on the reactants.

6. An increase resulting from the breakage of the full ring to yield an intermediate of length j (positive term). The rate γ represents the breakage of *two* interactions in the full ring to yield the smaller intermediate (see section 1):

$$M_j^{(6)} = n \cdot \gamma \cdot X_n.$$

This term covers both class 7 and class 8 dissociation reactions, since both cases have the same net effect on the products (see section 2.1.2).

Full Ring. The dynamics of the concentration of the full ring (X_n) can be described in terms of two processes: the binding of complementary intermediates to form the ring, and the dissociation of the ring due to the simultaneous breaking of two interactions. These terms are denoted $M_n^{(1)}$ and $M_n^{(2)}$, respectively:

$$\frac{dX_n}{dt} = M_n^{(1)} + M_n^{(2)}. \quad (7)$$

1. The first term, $M_n^{(1)}$, models the process of the binding of complementary intermediates to form the ring, and hence consists of positive terms:

$$M_n^{(1)} = \alpha \sum_{l=1}^{\lfloor \frac{n}{2} \rfloor} \left(\frac{1}{1 + \delta_{j,n/2}} \right) X_l \cdot X_{n-l},$$

where the floor function $\lfloor x \rfloor$ returns the largest integer smaller than x and arises to ensure correct counting for both odd and even rings. The $\delta_{j,n/2}$ term is the standard Kronecker delta and is included to represent the difference between type 3 and type 4 association reactions (see section 2.1.2): when $j \neq n/2$, the reaction occurs at rate $\alpha \cdot X_l \cdot X_{n-l}$, but when $j = n/2$ the reaction occurs with rate $(\alpha/2) \cdot X_{n/2}^2$.

2. The second term, $M_n^{(2)}$, models the process of the breakage of the ring by dissociation of two interactions. There are $\binom{n}{2}$ ways of choosing any two of the n interactions in the full ring to dissociate with rate γ , yielding the following expression:

$$M_n^{(2)} = -\binom{n}{2} \cdot \gamma \cdot X_n.$$

Note that this covers the net effect of all class 7 and class 8 dissociation reactions described in section 2.1.2.

2.1.4 Example: homomeric three-ring

The full set of differential equations for the homomeric three-ring can therefore be given as follows:

$$\begin{aligned} \frac{dX_1}{dt} &= 2\beta X_2 - 2\alpha X_1^2 - \alpha X_1 X_2 + 3\gamma X_3 \\ \frac{dX_2}{dt} &= \alpha X_1^2 - \beta X_2 - \alpha X_2 X_1 + 3\gamma X_3 \\ \frac{dX_3}{dt} &= \alpha X_1 X_2 - 3\gamma X_3. \end{aligned} \tag{8}$$

2.2 Assembly of heteromeric rings

2.2.1 Notation

We now consider a ring structure with n distinct subunits, which we label x_0, x_1, \dots, x_{n-1} —in this case, the subscripts refer to the *identity* of the subunit, not the size of a sub-complex as in the heteromeric case. Because the subunits are distinct, and each subunit can bind only to its neighbors, the ring is *heteromeric* (see Fig. 1 in the main text). In our notation, the indices for the subunits represent the equivalence classes $[0], [1], \dots, [n-1]$ of the integer modular arithmetic group \mathbb{Z}_n , where n is the ring length. That is, counting, addition, and subtraction on these indices in the discussion below should be understood to occur modulo n . In classifying the various sub-complexes of the full ring we adopt the convention that we enumerate the subunits in increasing order, identifying the complex by the index of the subunit from which we begin counting; so for example, in the case of a ring of length five, the trimer containing the subunits x_4, x_0 and x_1 can be identified as a molecule containing three subunits starting from subunit x_4 , the “first” subunit of the complex. We refer to the subunits that come first in counting order as occupying the “left” end of the complex, while those that come last in counting order are on the “right” end: in the trimer described above, subunit x_4 is on the left end, while subunit x_1 is on the right end.

Let $x_{i,j}$ denote the ring assembly intermediate starting from subunit i and containing j total subunits. For example, in the case of a five-membered ring, $x_{2,1}$ denotes the monomer x_2 , while $x_{3,4}$ denotes the tetramer starting at subunit x_3 and including subunits x_4, x_0 , and x_1 . For any given sub-complex of size $j < n$ there are n distinct sub-complexes of that size. An issue arises in

our notation due to the fact that any of the variables $x_{i,n}$ could be used to denote the concentration of the full ring, since they all refer to the same species. For simplicity, we use the variable “ x_n ” to denote the full ring, with $x_n \equiv x_{i,n} \forall i$. Since the individual species in the ring x_i have indices ranging from 0 to $n - 1$, the variable x_n refers unambiguously to the full ring.

As in the homomeric case, we use the variable “ $x_{i,j}$ ” to represent both the concentration of a given species and as a label for the species itself; the meaning of any given $x_{i,j}$ will be clear from the context.

Since the binding interfaces between subunits are distinct and may vary in strength, we can no longer apply a single off rate β to all dissociation reactions as we did in the homomeric case. We denote the off rate between two subunits x_i and x_{i+1} as β_i ; that is, off rates are labeled from the left side. Similarly, when describing the rate of ring breakage, we can no longer apply a single rate γ . Since the ring breakage rate is determined by the strength of the interactions at the two points of breakage, in the heteromeric case we must denote the two interactions that break. As in equation 4, we use the notation $\gamma_{i,j}$ to denote the rate of ring breakage occurring at two junctions: between subunits i and $i + 1$, and j and $j + 1$. This rate $\gamma_{i,j}$ is determined by the strengths of the interactions between these subunits (dissociation constants $K_D(i)$ and $K_D(j)$). See equation 4 in section 1 for a description of how to calculate $\gamma_{i,j}$ from these affinities.

The set of chemical reactions possible with heteromeric rings is similar to those described for homomeric rings in section 2.1.2. The difference here is that no reaction can ever consume or produce two copies of the same molecule since no intermediate in this system can ever react with another copy of itself in a productive way. We can thus ignore class 2, 4, 6 and 8 reactions as defined in section 2.1.2.

2.2.2 Ordinary differential equations

Assembly Intermediates. In analogy to equation 5 for the homomeric case, the differential equations describing the change in concentration of an assembly intermediate $x_{i,j}$ have the following form:

$$\frac{dx_{i,j}}{dt} = T_{i,j}^{(1)} + T_{i,j}^{(2)} + T_{i,j}^{(3)} + T_{i,j}^{(4)} + T_{i,j}^{(5)} + T_{i,j}^{(6)} \quad 1 \leq j < n, \ 0 \leq i < n, \quad (9)$$

where the $T_{i,j}^{(k)}$ represent groups of terms describing the exchange of mass between intermediates resulting from the same six binding and unbinding processes described above for the homomeric case. In the heteromeric case, these terms have the the following structure:

1. An increase in $x_{i,j}$ resulting from the dissociation of $x_{i,j}$ from either end of higher-order intermediates, yielding a term consisting of two sums:

$$T_{i,j}^{(1)} = \beta_{i+j-1} \sum_{l=j+1}^{n-1} x_{i,l} + \beta_{i-1} \sum_{l=1}^{n-j-1} x_{i-l,j+l}.$$

2. An increase resulting from smaller intermediates binding together to form $x_{i,j}$ (positive

terms):

$$T_{i,j}^{(2)} = \alpha \sum_{l=1}^{j-1} x_{i,l} \cdot x_{i+l,j-l}.$$

3. A decrease resulting from the binding of $x_{i,j}$ to complementary intermediates to form higher-order complexes (negative terms), but not the full ring. Since binding can occur at either of the two ends of $x_{i,j}$, this term consists of two sums:

$$T_{i,j}^{(3)} = -\alpha \cdot x_{i,j} \left(\sum_{l=1}^{n-j-1} x_{i+j,l} + \sum_{l=1}^{n-j-1} x_{i-l,l} \right).$$

4. A decrease resulting from the breakage of one of the $j - 1$ interactions within the complex $x_{i,j}$ (negative term):

$$T_{i,j}^{(4)} = -x_{i,j} \sum_{k=i}^{i+j-2} \beta_k,$$

where the sum is limited by $k = i + j - 2$ since the “final” interaction that can be broken in the complex occurs between subunits $i + j - 2$ and $i + j - 1$ and as such is indexed as β_{i+j-2} in our notation.

5. A decrease resulting from the binding of $x_{i,j}$ to its complementary intermediate to form the full ring (negative term):

$$T_{i,j}^{(5)} = -\alpha \cdot x_{i,j} \cdot x_{i+j,n-j}.$$

6. An increase resulting from the breakage of the full ring to yield the intermediate $x_{i,j}$ (positive term). There is only one pair of interactions that can break to produce the intermediate $x_{i,j}$: the interaction joining subunit x_{i-1} to subunit x_i (at the “left” end of $x_{i,j}$) and the interaction joining subunit x_{i+j-1} to subunit x_{i+j} (at the “right” end). This compound rate is denoted $\gamma_{i-1,i+j-1}$, yielding the following expression:

$$T_j^{(6)} = \gamma_{i-1,i+j-1} \cdot x_n.$$

Full Ring. As in the homomeric case, there is a single equation for the full ring, and it incorporates two processes: the binding of complementary intermediates to form the ring, and the breakage of the ring due to the dissociation of two interactions. These terms are denoted $T_n^{(1)}$ and $T_n^{(2)}$, respectively:

$$\frac{dx_n}{dt} = T_n^{(1)} + T_n^{(2)}. \quad (10)$$

1. The first term, $T_n^{(1)}$, models the binding of complementary intermediates to form the ring,

and hence consists of positive terms:

$$T_n^{(1)} = \alpha \sum_{j=1}^{\lfloor \frac{n}{2} \rfloor} \left(\sum_{i=0}^{n-1} \left(\frac{1}{1 + \delta_{j,n/2}} \right) x_{i,j} \cdot x_{i+j,n-j} \right).$$

The Kronecker delta ($\delta_{j,n/2}$) arises to compensate for the fact that the inner sum above *double-counts* the distinct reactions that can occur between intermediates of length $n/2$.

2. The second term, $T_n^{(2)}$, models the process of the breakage of the ring by dissociation of two interactions. Since dissociation can occur at any two interactions, we must sum over all of the unique $\gamma_{i,j}$ to get the total rate of ring breakage. Note that because the breakage of the two interactions is understood to occur simultaneously the order of interaction breakage does not matter: the rate $\gamma_{i,j}$ is therefore equivalent to $\gamma_{j,i}$ and only one of these permutations is counted when taking the sum.

$$T_n^{(2)} = -x_n \sum_{i=0}^{n-2} \left(\sum_{j=i+1}^{n-1} \gamma_{i,j} \right).$$

2.2.3 Example: heteromeric three-ring

As a specific example, the full set of equations for the change in concentration of $x_{i,j}$ for a three-membered ring is therefore:

$$\begin{aligned} \frac{dx_{0,1}}{dt} &= \beta_0 x_{0,2} + \beta_2 x_{2,2} - \alpha x_{0,1} (x_{1,1} + x_{2,1}) - \alpha x_{0,1} x_{1,2} + \gamma_{2,0} x_3 \\ \frac{dx_{1,1}}{dt} &= \beta_1 x_{1,2} + \beta_0 x_{0,2} - \alpha x_{1,1} (x_{2,1} + x_{0,1}) - \alpha x_{1,1} x_{2,2} + \gamma_{0,1} x_3 \\ \frac{dx_{2,1}}{dt} &= \beta_2 x_{2,2} + \beta_1 x_{1,2} - \alpha x_{2,1} (x_{0,1} + x_{1,1}) - \alpha x_{2,1} x_{0,2} + \gamma_{1,2} x_3 \\ \frac{dx_{0,2}}{dt} &= \alpha x_{0,1} x_{1,1} - \beta_0 x_{0,2} - \alpha x_{0,2} x_{2,1} + \gamma_{2,1} x_3 \\ \frac{dx_{1,2}}{dt} &= \alpha x_{1,1} x_{2,1} - \beta_1 x_{1,2} - \alpha x_{1,2} x_{0,1} + \gamma_{0,2} x_3 \\ \frac{dx_{2,2}}{dt} &= \alpha x_{2,1} x_{0,1} - \beta_2 x_{2,2} - \alpha x_{2,2} x_{1,1} + \gamma_{1,0} x_3 \\ \frac{dx_3}{dt} &= \alpha (x_{0,1} x_{1,2} + x_{1,1} x_{2,2} + x_{2,1} x_{0,2}) + x_3 (\gamma_{0,1} + \gamma_{0,2} + \gamma_{1,2}). \end{aligned}$$

2.3 Homomeric ring assembly arises as a special case of heteromeric ring assembly

Using the equations shown in Sections 2.1.3 and 2.2.2, we can show that the assembly of homomeric rings arises as a special case of heteromeric ring assembly where the affinities and initial subunit concentrations are all equal.

Assembly Intermediates. We start by considering heteromeric ring assembly. The concentration of all heteromeric intermediates of length $j < n$, denoted as X'_j , is simply the sum of the concentrations of the various $x_{i,j}$ intermediates of that length:

$$X'_j \equiv \sum_{i=0}^{n-1} x_{i,j} \quad 1 \leq j < n.$$

Consider the case where the initial concentrations of each subunit are equal, and where affinities between subunits are uniform—that is, $\beta_i = \beta \forall i$ and thus $\gamma_{i,j} = \gamma \forall (i,j)$. In this situation, at time $t = 0$ terms 1-6 in equation 9 will be equal for any two intermediates of length $j < n$; that is, we will have $T_{i,j}^x = T_{k,j}^x$ for any term x in equation 9 and any pair of species $x_{i,j}$ and $x_{k,j}$. At $t = 0$ we thus have that the differential equations for any individual intermediate of length j will be equal to the differential equations for all the other intermediates of length j . Since there is no source of symmetry breaking in this case, the differential equations will remain identical throughout the time evolution of the system, and so we have $x_{i,j} = x_{k,j}$ at all times. This yields:

$$X'_j = \sum_{k=0}^{n-1} x_{k,j} = n \cdot x_{i,j} \quad 1 \leq j < n$$

$$x_{i,j} = \frac{X'_j}{n}, \tag{11}$$

where we have arbitrarily chosen an $x_{i,j}$ to represent all intermediates of length j since they are all equal. By the chain rule, the differential equation for the change in concentration of X'_j is the sum of the equations for all $x_{i,j}$:

$$\frac{dX'_j}{dt} = \sum_{k=0}^{n-1} \frac{dx_{k,j}}{dt} = n \cdot \frac{dx_{i,j}}{dt} \quad 1 \leq j < n. \tag{12}$$

As in the homomeric case (equation 5) and heteromeric case (equation 9), the differential equation $\frac{dX'_j}{dt}$ consists of terms corresponding to the six processes of binding and unbinding:

$$\frac{dX'_j}{dt} = S_j^{(1)} + S_j^{(2)} + S_j^{(3)} + S_j^{(4)} + S_j^{(5)} + S_j^{(6)}.$$

We start from the heteromeric equations defining the terms of $\frac{dx_{i,j}}{dt}$, denoting the heteromeric on rate as α_t to distinguish it from the homomeric on rate, denoted α_m . Using the relations given by equations 11 and 12, we can write the equations for $\frac{dX'_j}{dt}$ as follows:

1. We can write $S_j^{(1)}$ by starting with the heteromeric term $T_{i,j}^{(1)}$, substituting $\frac{X'_j}{n}$ for any $x_{i,j}$ (by equation 11), substituting β for any β_i , and multiplying by the factor n (by equation

12), to get :

$$\begin{aligned}
S_j^{(1)} &= n \cdot \beta \left(\sum_{l=j+1}^{n-1} \frac{X'_l}{n} + \sum_{l=1}^{n-j-1} \frac{X'_{j+l}}{n} \right) \\
&= \beta \left(\sum_{l=j+1}^{n-1} X'_l + \sum_{l=1}^{n-j-1} X'_{j+l} \right) \\
&= \beta \left((X'_{j+1} + X'_{j+2} + \dots + X'_{n-1}) + (X'_{j+1} + X'_{j+2} + \dots + X'_{n-1}) \right).
\end{aligned}$$

As the final line above demonstrates, the two summations in the expression sum over the same set of terms and hence can be replaced by a single sum multiplied by a factor of two. We thus find that $S_j^{(1)}$ is equivalent to the term $M_j^{(1)}$:

$$S_j^{(1)} = 2 \cdot \beta \sum_{l=j+1}^{n-1} X'_l = M_j^{(1)}.$$

2. We repeat the above procedure, starting with the terms in $T_{i,j}^{(2)}$ for the heteromeric case and substituting $\frac{X'_j}{n}$ for any $x_{i,j}$ and multiplying by the factor n :

$$\begin{aligned}
S_j^{(2)} &= n \cdot \alpha_t \sum_{l=1}^{j-1} \frac{X'_l}{n} \cdot \frac{X'_{j-l}}{n} \\
&= \frac{\alpha_t}{n} \sum_{l=1}^{j-1} X'_l \cdot X'_{j-l}.
\end{aligned}$$

Note that, if we introduce the relationship $\alpha_m = \frac{\alpha_t}{n}$, this yields:

$$S_j^{(2)} = \alpha_m \sum_{l=1}^{j-1} X'_l \cdot X'_{j-l} = M_j^{(2)}.$$

We thus have that the heteromeric term $S_j^{(2)} = M_j^{(2)}$ when we set the homomeric association rate α_m to equal the heteromeric association rate α_t , divided by the ring length n .

3. We calculate $S_j^{(3)}$ from $T_{i,j}^{(3)}$:

$$\begin{aligned} S_j^{(3)} &= -n \cdot \alpha_t \cdot \frac{X'_j}{n} \left(\sum_{l=1}^{n-j-1} \frac{X'_l}{n} + \sum_{l=1}^{n-j-1} \frac{X'_l}{n} \right) \\ &= -2 \cdot \frac{\alpha_t}{n} \cdot X'_j \sum_{l=1}^{n-j-1} X'_l \\ &= -2 \cdot \alpha_m \cdot X'_j \sum_{l=1}^{n-j-1} X'_l = M_j^{(3)}, \end{aligned}$$

where we have again used the substitution $\alpha_m = \frac{\alpha_t}{n}$.

4. To calculate $S_j^{(4)}$, we begin with $T_{i,j}^{(4)}$, choosing $i = 0$ for simplicity in the following sums:

$$\begin{aligned} S_j^{(4)} &= -n \cdot \frac{X'_j}{n} \sum_{k=0}^{j-2} \beta_k \\ &= -X'_j \cdot (\beta_0 + \beta_1 + \dots + \beta_{j-2}). \end{aligned}$$

Since all $\beta_i = \beta$, and since there are $j - 1$ terms in the sum, we can substitute to yield an expression identical to the homomeric equation:

$$S_j^{(4)} = -\beta \cdot X'_j \cdot (j - 1) = M_j^{(4)}.$$

5. We repeat the procedure starting with $T_{i,j}^{(5)}$, again noting the transformation $\alpha_m = \frac{\alpha_t}{n}$:

$$\begin{aligned} S_j^{(5)} &= -n \cdot \alpha_t \cdot \frac{X'_j}{n} \cdot \frac{X'_{n-j}}{n} \\ &= -\alpha_m \cdot X'_j \cdot X'_{n-j} = M_j^{(5)}. \end{aligned}$$

6. We repeat the procedure for $T_{i,j}^{(6)}$, which describes the process of ring breakage. We denote the concentration of the full heteromeric ring in terms of the combined heteromeric intermediates X'_j as X'_n . Since in the heteromeric case there is only instance of the full ring x_n , $X'_n = x_n$. Due to the assumed uniformity of affinities, $\gamma_{i,j} = \gamma$:

$$S_j^{(6)} = n \cdot \gamma \cdot x_n = M_j^{(6)}.$$

Full Ring. To determine the equation for the full heteromeric ring X'_n in terms of the combined heteromeric intermediates X'_j , we start with the equation for the full heteromeric ring as given in equation 10 and make the appropriate substitutions. As in the heteromeric and homomeric cases, the equation consists of two terms describing the processes of ring formation and ring breakage:

$$\frac{dX'_n}{dt} = S_n^{(1)} + S_n^{(2)}.$$

1. We start with the first term, $S_n^{(1)}$, describing the process of ring formation. As in the derivation for the assembly intermediates, we substitute $\frac{X'_j}{n}$ for any $x_{i,j}$; in this case, however, we do not multiply the entire term by n , since $X'_n = x_n$ in contrast to the assembly intermediates, where $X'_{i,j} = n \cdot x_{i,j}$.

$$S_n^{(1)} = \alpha_t \sum_{j=1}^{\lfloor \frac{n}{2} \rfloor} \left(\sum_{i=0}^{n-1} \left(\frac{1}{1 + \delta_{j,n/2}} \right) \frac{X'_j}{n} \cdot \frac{X'_{n-j}}{n} \right).$$

We note that the terms in the inner sum do not depend on i , and thus this sum represents multiplication by a factor of n . Remembering $\alpha_m = \frac{\alpha_t}{n}$, we have:

$$\begin{aligned} S_n^{(1)} &= n \cdot \alpha_t \sum_{j=1}^{\lfloor \frac{n}{2} \rfloor} \left(\frac{1}{1 + \delta_{j,n/2}} \right) \frac{X'_j}{n} \cdot \frac{X'_{n-j}}{n} \\ &= \alpha_m \sum_{j=1}^{\lfloor \frac{n}{2} \rfloor} \left(\frac{1}{1 + \delta_{j,n/2}} \right) X'_j \cdot X'_{n-j} = M_n^{(1)}. \end{aligned}$$

2. We repeat the procedure for the second term, $S_n^{(2)}$:

$$\begin{aligned} S_n^{(2)} &= -X'_n \sum_{i=0}^{n-2} \left(\sum_{j=i+1}^{n-1} \gamma_{i,j} \right) \\ &= -X'_n \sum_{i=0}^{n-2} (\gamma_{i,i+1} + \gamma_{i,i+2} + \dots + \gamma_{i,n-1}). \end{aligned}$$

There are $n - 1 - i$ terms in the inner sum, and since $\gamma_{i,j} = \gamma$ due to interaction uniformity, this yields:

$$\begin{aligned} S_n^{(2)} &= -X'_n \sum_{i=0}^{n-2} \gamma(n - 1 - i) \\ &= -\gamma \cdot X'_n ((n - 1) + (n - 2) + \dots + (n - 1 - (n - 3)) + (n - 1 - (n - 2))) \\ &= -\gamma \cdot X'_n ((n - 1) + (n - 2) + \dots + 2 + 1)) \\ &= -\gamma \cdot X'_n \sum_{k=1}^{n-1} k. \end{aligned}$$

We can rewrite this summation, $\sum_{k=1}^{n-1} k$, in terms of the well-known sum of the first n

natural numbers: $\sum_{k=1}^{n-1} k = (\sum_{k=1}^n k) - n = n(n+1)/2 - n$. This yields:

$$\begin{aligned} S_n^{(2)} &= -\left(\frac{n^2 + n - 2n}{2}\right) \cdot \gamma \cdot X'_n \\ &= -\frac{n(n-1)}{2} \cdot \gamma \cdot X'_n \\ &= -\binom{n}{2} \cdot \gamma \cdot X'_n = M_n^{(2)}. \end{aligned}$$

Summary. As described above, we can show that every $S_j^{(x)}$ is equivalent to the corresponding $M_j^{(x)}$ term. This indicates that we can describe the time evolution of the size-class variables X'_j in the heteromeric case with the same exact equations governing the dynamics of the homomeric X_j concentrations. The assembly dynamics of any given heteromeric ring in which the following conditions hold:

1. The off rates β_i and thus the ring breakage rates $\gamma_{i,j}$ are equal for all subunits in the heteromeric case
2. For any given length j , the initial concentrations of the complexes $x_{i,j}$ are equal for all i

can thus be represented by a homomeric system with on rate $\alpha_m = \frac{\alpha t}{n}$. The rescaling of the association rate arises naturally in this case: given some $k + l \leq n$, a given homomeric X_k molecule can react with any X_l molecule. Any heteromeric $x_{i,k}$ molecule, however, can only react with one of the n members of X'_l . Homomeric molecules thus have n more binding options than heteromeric ones, and dividing the association rate by n compensates for this fact.

2.4 Modeling synthesis and degradation of rings

The models described in sections 2.1.3 and 2.2.2 are meant to capture the dynamics of ring assembly when a closed pool of monomers is allowed to spontaneously assemble. This situation mimics the conditions of many *in vitro* assembly experiments (e.g., [5]), or cases in which assembly is characterized by rapid activation of monomers (e.g. the apoptosome [6]). As discussed in the main text, however, most cells do not synthesize a large number of monomers and only then allow them to interact—assembly in the majority of cases is likely to be characterized by a steady-state process in which the cell constantly produces new monomers to compensate for continuous decreases in the concentration of both assembly intermediates and the full ring due to cell growth or active protein degradation.

We must thus understand whether the dynamic deadlock we have identified (i.e. the plateau in Fig. 2A in the main text) has any bearing on the assembly of rings in the presence of synthesis and degradation. These two processes could each have a role in facilitating assembly: the degradation of persisting intermediates alters or removes the incompatible complexes responsible for deadlock, while the synthesis of new monomers allows any persisting intermediates to find compatible binding partners and complete the process of assembly.

To explore how subunit synthesis and degradation of complexes would affect ring assembly dynamics, we designed additional models incorporating these two processes. Existing experimental evidence suggests that many distinct “degradation” scenarios are possible, depending on the particular cell type in question. We consider two specific cases.

In the first model, which we denote “subunit deletion” or simply “model A,” we assume that each individual *subunit* has a certain likelihood of being degraded, irrespective of its membership in a protein complex. When a subunit is deleted from a complex, the other subunits are left intact. In the context of a system in which proteins are post-translationally tagged and targeted for degradation by proteases, this implies that (1) the probability of a protein being tagged is insensitive to binding context, and (2) the degradation machinery is capable of removing individual subunits from complexes. Experimental work on the bacterial protease ClpXP suggests that it can extract and degrade a tagged subunit from a complex in precisely this fashion, leaving other bound (but untagged) subunits intact [7]; this capability of ClpXP has also been exploited for use experimentally as “molecular tweezers” [8].

In the second model, denoted “whole-complex degradation” or “model B,” we instead assume that each *complex* has a certain likelihood of being degraded—this includes both monomers and fully assembled rings. The simplest physical interpretation of this model is one of continuous dilution: for example, if the cell grows continuously in volume, the concentrations of all complexes, monomers and rings alike, will be reduced. In the context of active protein degradation processes this model assumes that (1) the propensity for degradation is a property of complexes rather than subunits, and hence that all complexes have an equal likelihood of being tagged regardless of size, and (2) that the degradation machinery consumes the entire complex in one interaction. In contrast to ClpXP, the bacterial protease ClpAP appears to use this alternative mechanism, degrading whole heterodimeric complexes even in cases where only subunit has been tagged [7].

As described above, there is some experimental evidence for the degradation processes described in model A and model B. These two models, however, should not be considered comprehensive, nor mutually exclusive—the dynamics of degradation processes in cells are likely to depend on complex size, specific subunit affinity for targeting enzymes, the number and type of degradative post-translational modifications, and many other factors in a highly complex manner. In this work we consider these two cases separately for the sake of simplicity; we leave exploration of the effects of more complicated degradation processes to future work.

In the following sections we describe the ordinary differential equations used to model homomeric and heteromeric ring assembly in the presence of model A- or model B-type degradation. In both cases, production of monomers occurs at a rate Q that is considered to be constant in time, and all first-order degradation processes occur with a rate δ .

2.4.1 Homomeric model A: subunit deletion

Assembly Intermediates with Lengths between 2 and $n - 1$. As shown in equation 5, the ordinary differential equations for the change in concentration of any homomeric ring-assembly intermediate $X_j, 1 \leq j < n$ can be written as the sum of six groups of terms, $M_j^{(1)} + \dots + M_j^{(6)}$, describing six specific binding and unbinding processes. To incorporate subunit deletion into our

equations describing the concentrations of assembly intermediates we add additional terms to each $\frac{dX_j}{dt}$. For intermediates of length $2 \leq j < n - 1$ (we deal with monomers in a separate section below), these terms take the form:

$$\frac{dX_j}{dt} = M_j^{(1)} + \dots + M_j^{(6)} + M_j^{(A1)} + M_j^{(A2)}. \quad 2 \leq j < n - 1 \quad (13)$$

The physical processes modeled by these two additional terms are as follows:

1. A decrease in X_j resulting from deletion of a monomer from the complex (negative term). In model A, each monomer has the same independent chance of being degraded; since there are j monomers in X_j , the deletion rate δ is multiplied by j :

$$M_j^{(A1)} = -j \cdot \delta \cdot X_j. \quad (14)$$

2. An increase in X_j resulting from deletion of a monomer from a larger intermediate X_k of size $j < k < n$, yielding X_j (positive terms). Because there are exactly two monomers that can be deleted from X_k to yield a smaller intermediate of length j , the degradation rate δ is multiplied by two:

$$M_j^{(A2)} = 2 \cdot \delta \cdot \sum_{k=j+1}^{n-1} X_k.$$

Monomers. The equation for monomers X_1 is identical to equation 13 for intermediates smaller than length $n - 1$, with the addition of a single additional term, Q , denoting the rate of synthesis of new monomers:

$$\frac{dX_1}{dt} = M_1^{(1)} + \dots + M_1^{(6)} + M_1^{(A1)} + M_1^{(A2)} + Q.$$

Assembly Intermediates of Length $n - 1$. The equation for assembly intermediates X_{n-1} is similar to the equations for smaller intermediates; however, since the only larger intermediate (from which a monomer can be deleted to yield X_{n-1}) is the full ring X_n , the equation has the following structure:

$$\frac{dX_{n-1}}{dt} = M_{n-1}^{(1)} + \dots + M_{n-1}^{(6)} + M_{n-1}^{(A1)} + M_{n-1}^{(A3)}.$$

In this expression $M_{n-1}^{(A1)}$ is as in equation 14 with $j = n - 1$, and $M_{n-1}^{(A3)}$ is as follows:

$$M_{n-1}^{(A3)} = n \cdot \delta \cdot X_n.$$

The factor n is applied to the rate δ due to the fact that there are n subunits that can be deleted from the full ring X_n to yield the intermediate X_{n-1} .

Full Ring. The formation of the full ring X_n with the addition of degradation includes the two processes described for equation 7, with the addition of a term to describe the degradation of the

full ring:

$$\frac{dX_n}{dt} = M_n^{(1)} + M_n^{(2)} + M_n^{(A4)}.$$

Because there are n subunits that can be deleted from the full ring, each with rate δ , the term $M_n^{(A4)}$ has the following form:

$$M_n^{(A4)} = -n \cdot \delta \cdot X_n.$$

Example: Homomeric Three-Ring. The full set of differential equations for the homomeric three-ring with subunit deletion is as follows:

$$\begin{aligned} \frac{dX_1}{dt} &= 2\beta X_2 - 2\alpha X_1^2 - \alpha X_1 X_2 + 3\gamma X_3 - \delta X_1 + 2\delta X_2 + Q \\ \frac{dX_2}{dt} &= \alpha X_1^2 - \beta X_2 - \alpha X_2 X_1 + 3\gamma X_3 - 2\delta X_2 + 3\delta X_3 \\ \frac{dX_3}{dt} &= \alpha X_1 X_2 - 3\gamma X_3 - 3\delta X_3. \end{aligned} \tag{15}$$

2.4.2 Heteromeric model A: subunit deletion

We now consider the case of heteromeric ring assembly with model A degradation, where individual subunits can be deleted individually from assembled complexes. Each subunit x_i may have a different rate of deletion, denoted δ_i ; the synthesis rate of monomer x_i is denoted Q_i . Though we describe here the form of the equations allowing for differences between the different δ_i and Q_i , in our simulations we considered only the case where the synthesis and degradation rates of all subunits are equal, that is $\delta_i = \delta$ and $Q_i = Q$.

Assembly Intermediates with Lengths between 2 and $n - 1$. As shown in equation 9, the ordinary differential equations for the change in concentration of any heteromeric ring-assembly intermediate $x_{i,j}$, $1 \leq j < n$, $0 \leq i < n$ can be written as the sum of six groups of terms, $T_{i,j}^{(1)} + \dots + T_{i,j}^{(6)}$, describing six specific binding and unbinding processes. As in the case of homomeric rings, to incorporate subunit deletion into our equations describing the concentrations of assembly intermediates we add additional terms to each $\frac{dx_{i,j}}{dt}$. For intermediates of length $2 \leq j < n - 1$ (as for the homomeric case we deal with monomers in a separate section below), these terms take the form:

$$\frac{dx_{i,j}}{dt} = T_{i,j}^{(1)} + \dots + T_{i,j}^{(6)} + T_{i,j}^{(A1)} + T_{i,j}^{(A2)} \quad 2 \leq j < n - 1, 0 \leq i < n. \tag{16}$$

The physical processes modeled by these two additional terms are as follows:

1. A decrease in $x_{i,j}$ resulting from deletion of a monomer from the complex (negative terms). Since any of the subunits within $x_{i,j}$ can be deleted, we sum over the deletion rates for all j

members of $x_{i,j}$:

$$T_{i,j}^{(A1)} = -x_{i,j} \sum_{k=0}^{j-1} \delta_{i+k}. \quad (17)$$

2. An increase in $x_{i,j}$ resulting from deletion of a monomer from a larger intermediate of size $j < k < n$, yielding $x_{i,j}$ (positive terms). For deletion of a single subunit to result in the production of $x_{i,j}$, $x_{i,j}$ must be on one of the two “ends” of the larger complex. In the complexes where $x_{i,j}$ is on the “left end” (using the notation described in section 2.2.1), $x_{i,j}$ will be produced when the subunit that links $x_{i,j}$ to the complex is deleted; this is subunit x_{i+j} . In complexes where $x_{i,j}$ is on the “right end”, subunit x_{i-1} must be deleted in order to obtain $x_{i,j}$. We thus obtain:

$$T_{i,j}^{(A2)} = \delta_{i+j} \sum_{k=j+1}^{n-1} x_{i,k} + \delta_{i-1} \sum_{k=j+1}^{n-1} x_{i+j-k,k},$$

where the first sum considers all the left cases, and the second all the right cases.

Monomers. The equation for monomers $x_{i,1}$ is identical to equation 16 for intermediates smaller than length $n - 1$, with the addition of a single additional term, Q_i , denoting the rate of synthesis of monomer i :

$$\frac{dx_{i,1}}{dt} = T_{i,1}^{(1)} + \dots + T_{i,1}^{(6)} + T_{i,1}^{(A1)} + T_{i,1}^{(A2)} + Q_i.$$

Assembly Intermediates of Length $n - 1$. The equation for assembly intermediates $x_{i,n-1}$ is similar to the equations for smaller intermediates; however, since the only larger intermediate from which a monomer can be deleted to yield $x_{i,n-1}$ is the full ring x_n , we have:

$$\frac{dx_{i,n-1}}{dt} = T_{i,n-1}^{(1)} + \dots + T_{i,n-1}^{(6)} + T_{i,n-1}^{(A1)} + T_{i,n-1}^{(A3)}.$$

In this expression $T_{i,n-1}^{(A1)}$ is as in equation 17 with $j = n - 1$. $T_{i,n-1}^{(A3)}$ describes the rate of formation of $x_{i,n-1}$ by deletion of a subunit from the full ring; since the only monomer that can be deleted from the full ring to result in $x_{i,n-1}$ is the monomer x_{i-1} , $T_{i,n-1}^{(A3)}$ is written:

$$T_{i,n-1}^{(A3)} = \delta_{i-1} \cdot x_n.$$

Full Ring. The equation for formation of the full ring x_n includes the two assembly processes described in equation 10, with the addition of a term to describe the degradation of the full ring:

$$\frac{dx_n}{dt} = T_n^{(1)} + T_n^{(2)} + T_n^{(A4)}.$$

Since any of the subunits x_i can be deleted from the full ring, each with its own rate δ_i , the term has the following form:

$$T_n^{(A4)} = -x_n \sum_{i=0}^{n-1} \delta_i.$$

Example: Heteromeric Three-Ring. The full set of differential equations for the heteromeric three-ring with subunit deletion is as follows:

$$\begin{aligned} \frac{dx_{0,1}}{dt} &= \beta_0 x_{0,2} + \beta_2 x_{2,2} - \alpha x_{0,1}(x_{1,1} + x_{2,1}) - \alpha x_{0,1} x_{1,2} + \gamma_{2,0} x_3 \\ &\quad - \delta_0 x_{0,1} + \delta_1 x_{0,2} + \delta_2 x_{2,2} + Q_0 \\ \frac{dx_{1,1}}{dt} &= \beta_1 x_{1,2} + \beta_0 x_{0,2} - \alpha x_{1,1}(x_{2,1} + x_{0,1}) - \alpha x_{1,1} x_{2,2} + \gamma_{0,1} x_3 \\ &\quad - \delta_1 x_{1,1} + \delta_2 x_{1,2} + \delta_0 x_{0,2} + Q_1 \\ \frac{dx_{2,1}}{dt} &= \beta_2 x_{2,2} + \beta_1 x_{1,2} - \alpha x_{2,1}(x_{0,1} + x_{1,1}) - \alpha x_{2,1} x_{0,2} + \gamma_{1,2} x_3 \\ &\quad - \delta_2 x_{2,1} + \delta_0 x_{2,2} + \delta_1 x_{1,2} + Q_2 \\ \frac{dx_{0,2}}{dt} &= \alpha x_{0,1} x_{1,1} - \beta_0 x_{0,2} - \alpha x_{0,2} x_{2,1} + \gamma_{2,1} x_3 \\ &\quad - x_{0,2}(\delta_0 + \delta_1) + \delta_2 x_3 \\ \frac{dx_{1,2}}{dt} &= \alpha x_{1,1} x_{2,1} - \beta_1 x_{1,2} - \alpha x_{1,2} x_{0,1} + \gamma_{0,2} x_3 \\ &\quad - x_{1,2}(\delta_1 + \delta_2) + \delta_0 x_3 \\ \frac{dx_{2,2}}{dt} &= \alpha x_{2,1} x_{0,1} - \beta_2 x_{2,2} - \alpha x_{2,2} x_{1,1} + \gamma_{1,0} x_3 \\ &\quad - x_{2,2}(\delta_2 + \delta_0) + \delta_1 x_3 \\ \frac{dx_3}{dt} &= \alpha(x_{0,1} x_{1,2} + x_{1,1} x_{2,2} + x_{2,1} x_{0,2}) + x_3(\gamma_{0,1} + \gamma_{0,2} + \gamma_{1,2}) \\ &\quad - x_3(\delta_0 + \delta_1 + \delta_2). \end{aligned}$$

2.4.3 Homomeric model B: whole-complex degradation

Assembly Intermediates. As shown in equation 5, the ordinary differential equations for the change in concentration of any ring-assembly intermediate X_j , $1 \leq j < n$ can be written as the sum of six groups of terms, $M_j^{(1)} + \dots + M_j^{(6)}$, describing six specific binding and unbinding processes. To incorporate whole-complex degradation into our equations describing the concentrations of assembly intermediates we add an additional term, $M_j^{(B1)}$, to each $\frac{dX_j}{dt}$. The resulting equation for intermediates of length $2 \leq j < n$ (monomers are dealt with in a separate section below) is as follows:

$$\frac{dX_j}{dt} = M_j^{(1)} + \dots + M_j^{(6)} + M_j^{(B1)} \quad 2 \leq j < n.$$

Since in this model of degradation all complexes have an equal rate δ of being degraded, the degradation term $M_j^{(B1)}$ is simply:

$$M_j^{(B1)} = -\delta \cdot X_j \quad 1 \leq j \leq n. \quad (18)$$

Monomers. In our equation for monomers X_1 , the degradation term $M_1^{(B1)}$ is as described above in equation 18, but we also account for the synthesis of new monomers by incorporating the synthesis rate Q :

$$\frac{dX_1}{dt} = M_1^{(1)} + \dots + M_1^{(6)} + M_1^{(B1)} + Q.$$

Full Ring. The equation for the formation of the full ring X_n consists of the same two terms as described for equation 7, with the addition of the degradation term $M_n^{(B1)}$:

$$\frac{dX_n}{dt} = M_n^{(1)} + M_n^{(2)} + M_n^{(B1)}. \quad (19)$$

$M_n^{(B1)}$ takes the same form as equation 18:

$$M_n^{(B1)} = -\delta \cdot X_n.$$

Example: Homomeric Three-Ring. The full set of differential equations for the homomeric three-ring with whole-complex degradation is as follows:

$$\begin{aligned} \frac{dX_1}{dt} &= 2\beta X_2 - 2\alpha X_1^2 - \alpha X_1 X_2 + 3\gamma X_3 - \delta X_1 + Q \\ \frac{dX_2}{dt} &= \alpha X_1^2 - \beta X_2 - \alpha X_2 X_1 + 3\gamma X_3 - \delta X_2 \\ \frac{dX_3}{dt} &= \alpha X_1 X_2 - 3\gamma X_3 - \delta X_3. \end{aligned} \quad (20)$$

2.4.4 Heteromeric model B: whole-complex degradation

Because the whole-complex degradation model applies the same degradation rate to all complexes, the differential equations for heteromeric ring assembly incorporating this degradation model require nearly identical modifications as those for homomeric ring assembly, described in section 2.4.3. Note that in this case while there is a single degradation rate δ that applies to all complexes, but we allow for different rates of monomer synthesis—the synthesis rate for monomer x_i is denoted Q_i . As in the case of the subunit deletion model for heteromeric rings (section 2.4.2), in our simulations we only consider the case where these rates are all equal (i.e. $Q_i = Q \forall i$).

Assembly intermediates. As shown in equation 9, the ordinary differential equations for the change in concentration of any heteromeric ring-assembly intermediate $x_{i,j}$, $1 \leq j < n$, $0 \leq i < n$ can be written as the sum of six groups of terms, $T_{i,j}^{(1)} + \dots + T_{i,j}^{(6)}$, describing six specific binding and

unbinding processes. As in the case of homomeric rings, to incorporate whole-complex degradation into our equations describing the concentrations of assembly intermediates we add a single additional term $T_{i,j}^{(B1)}$ to each $\frac{dx_{i,j}}{dt}$. The resulting equation for intermediates of length $2 \leq j < n$ (we deal with monomers in a separate section below), is as follows:

$$\frac{dx_{i,j}}{dt} = T_{i,j}^{(1)} + \dots + T_{i,j}^{(6)} + T_{i,j}^{(B1)} \quad 2 \leq j < n, \ 0 \leq i < n. \quad (21)$$

Since in this model of degradation all complexes have an equal rate δ of being degraded, the degradation term $T_{i,j}^{(B1)}$ is simply

$$T_{i,j}^{(B1)} = -\delta \cdot x_{i,j} \quad 1 \leq j \leq n. \quad (22)$$

Monomers. In our equation for monomers $x_{i,1}$, the degradation term $T_{i,1}^{(B1)}$ is as described above in equation 22, but we also account for the synthesis of new monomers by incorporating the synthesis rate Q_i :

$$\frac{dx_{i,j}}{dt} = T_{i,1}^{(1)} + \dots + T_{i,1}^{(6)} + T_{i,1}^{(B1)} + Q_i.$$

Full Ring. The formation of the full ring x_n with the addition of whole-complex degradation includes the two processes described for equation 10, with the addition of the term $T_n^{(B1)}$:

$$\frac{dx_n}{dt} = T_n^{(1)} + T_n^{(2)} + T_n^{(B1)}.$$

Here $T_n^{(B1)}$ is as described in equation 22:

$$T_n^{(B1)} = -\delta \cdot x_n.$$

Example: Heteromeric Three-Ring. The full set of differential equations for the heteromeric three-ring with whole-complex degradation is as follows:

$$\begin{aligned}
\frac{dx_{0,1}}{dt} &= \beta_0 x_{0,2} + \beta_2 x_{2,2} - \alpha x_{0,1}(x_{1,1} + x_{2,1}) - \alpha x_{0,1} x_{1,2} + \gamma_{2,0} x_3 \\
&\quad - \delta x_{0,1} + Q_0 \\
\frac{dx_{1,1}}{dt} &= \beta_1 x_{1,2} + \beta_0 x_{0,2} - \alpha x_{1,1}(x_{2,1} + x_{0,1}) - \alpha x_{1,1} x_{2,2} + \gamma_{0,1} x_3 \\
&\quad - \delta x_{1,1} + Q_1 \\
\frac{dx_{2,1}}{dt} &= \beta_2 x_{2,2} + \beta_1 x_{1,2} - \alpha x_{2,1}(x_{0,1} + x_{1,1}) - \alpha x_{2,1} x_{0,2} + \gamma_{1,2} x_3 \\
&\quad - \delta x_{2,1} + Q_2 \\
\frac{dx_{0,2}}{dt} &= \alpha x_{0,1} x_{1,1} - \beta_0 x_{0,2} - \alpha x_{0,2} x_{2,1} + \gamma_{2,1} x_3 \\
&\quad - \delta x_{0,2} \\
\frac{dx_{1,2}}{dt} &= \alpha x_{1,1} x_{2,1} - \beta_1 x_{1,2} - \alpha x_{1,2} x_{0,1} + \gamma_{0,2} x_3 \\
&\quad - \delta x_{1,2} \\
\frac{dx_{2,2}}{dt} &= \alpha x_{2,1} x_{0,1} - \beta_2 x_{2,2} - \alpha x_{2,2} x_{1,1} + \gamma_{1,0} x_3 \\
&\quad - \delta x_{2,2} \\
\frac{dx_3}{dt} &= \alpha x_{0,1} x_{1,2} + \alpha x_{1,1} x_{2,2} + \alpha x_{2,1} x_{0,2} + x_3(\gamma_{0,1} + \gamma_{0,2} + \gamma_{1,2}) \\
&\quad - \delta x_3.
\end{aligned}$$

2.5 Assembly of Heteromeric Chains

In order to provide a contrast to the affinity configurations that optimize the assembly efficiency of rings (section 4.3.2), we considered a model of the assembly of heteromeric chains (whose optimization we consider in section 4.3.3). We derived equations for the assembly of chains in a manner exactly analogous to that for rings as described in section 2.2. Given the similarities between the ring and chain case, we will not describe ODEs for chains of arbitrary length. We will instead restrict our discussion to chains of length four; these chains have exactly the same number of interactions as three-membered rings, allowing for direct comparison between the two cases (e.g. Fig. 4A in the main text).

In the equations that follow, we denote the monomers that make up the four-chain c_1, c_2, c_3 , and c_4 . By analogy to the notation for rings, $c_{i,j}$ is used to denote the concentration of the species that starts at subunit c_i and is of length j . β_i denotes the off rate for the interaction connecting subunit c_i to c_{i+1} . As with rings, α denotes the on rate, δ the degradation rate, and Q the monomer synthesis rate.

Assembly of a Heteromeric Chain of Length Four

$$\begin{aligned}
\frac{dc_{1,1}}{dt} &= \beta_1(c_{1,2} + c_{1,3} + c_{1,4}) - \alpha c_{1,1}(c_{2,1} + c_{2,2} + c_{2,3}) \\
\frac{dc_{2,1}}{dt} &= \beta_2(c_{2,2} + c_{2,3}) + \beta_1 c_{1,2} - \alpha c_{2,1}(c_{3,1} + c_{3,2} + c_{1,1}) \\
\frac{dc_{3,1}}{dt} &= \beta_3 c_{3,2} + \beta_2(c_{1,3} + c_{2,2}) - \alpha c_{3,1}(c_{4,1} + c_{2,1} + c_{1,2}) \\
\frac{dc_{4,1}}{dt} &= \beta_3(c_{1,4} + c_{2,3} + c_{3,2}) - \alpha c_{4,1}(c_{3,1} + c_{2,2} + c_{1,3}) \\
\frac{dc_{1,2}}{dt} &= \beta_2(c_{1,3} + c_{1,4}) + \alpha c_{1,1} c_{2,1} - \alpha c_{1,2}(c_{3,1} + c_{3,2}) - \beta_1 c_{1,2} \\
\frac{dc_{2,2}}{dt} &= \beta_3 c_{2,3} + \beta_1 c_{1,3} + \alpha c_{2,1} c_{3,1} - \alpha c_{2,2}(c_{4,1} + c_{1,1}) - \beta_2 c_{2,2} \\
\frac{dc_{3,2}}{dt} &= \beta_2(c_{1,4} + c_{2,3}) + \alpha c_{3,1} c_{4,1} - \alpha c_{3,2}(c_{2,1} + c_{1,2}) - \beta_3 c_{3,2} \\
\frac{dc_{1,3}}{dt} &= \beta_3 c_{1,4} + \alpha(c_{1,1} c_{2,2} + c_{1,2} c_{3,1}) - \alpha c_{1,3} c_{4,1} - c_{1,3}(\beta_1 + \beta_2) \\
\frac{dc_{2,3}}{dt} &= \beta_1 c_{1,4} + \alpha(c_{2,1} c_{3,2} + c_{2,2} c_{4,1}) - \alpha c_{2,3} c_{1,1} - c_{2,3}(\beta_2 + \beta_3) \\
\frac{dc_{1,4}}{dt} &= \alpha(c_{1,1} c_{2,3} + c_{1,2} c_{3,2} + c_{1,3} c_{4,1}) - c_{1,4}(\beta_1 + \beta_2 + \beta_3)
\end{aligned}$$

Assembly of a Heteromeric Chain with Model A Degradation

$$\begin{aligned}
\frac{dc_{1,1}}{dt} &= \beta_1(c_{1,2} + c_{1,3} + c_{1,4}) - \alpha c_{1,1}(c_{2,1} + c_{2,2} + c_{2,3}) - \delta c_{1,1} + \delta(c_{1,2} + c_{1,3} + c_{1,4}) + Q \\
\frac{dc_{2,1}}{dt} &= \beta_2(c_{2,2} + c_{2,3}) + \beta_1 c_{1,2} - \alpha c_{2,1}(c_{3,1} + c_{3,2} + c_{1,1}) - \delta c_{2,1} + \delta(c_{2,2} + c_{2,3} + c_{1,2}) + Q \\
\frac{dc_{3,1}}{dt} &= \beta_3 c_{3,2} + \beta_2(c_{1,3} + c_{2,2}) - \alpha c_{3,1}(c_{4,1} + c_{2,1} + c_{1,2}) - \delta c_{3,1} + \delta(c_{3,2} + c_{1,3} + c_{2,2}) + Q \\
\frac{dc_{4,1}}{dt} &= \beta_3(c_{1,4} + c_{2,3} + c_{3,2}) - \alpha c_{4,1}(c_{3,1} + c_{2,2} + c_{1,3}) - \delta c_{4,1} + \delta(c_{1,4} + c_{2,3} + c_{3,2}) + Q \\
\frac{dc_{1,2}}{dt} &= \beta_2(c_{1,3} + c_{1,4}) + \alpha c_{1,1} c_{2,1} - \alpha c_{1,2}(c_{3,1} + c_{3,2}) - \beta_1 c_{1,2} - 2\delta c_{1,2} + \delta(c_{1,3} + c_{1,4}) \\
\frac{dc_{2,2}}{dt} &= \beta_3 c_{2,3} + \beta_1 c_{1,3} + \alpha c_{2,1} c_{3,1} - \alpha c_{2,2}(c_{4,1} + c_{1,1}) - \beta_2 c_{2,2} - 2\delta c_{2,2} + \delta(c_{2,3} + c_{1,3}) \\
\frac{dc_{3,2}}{dt} &= \beta_2(c_{1,4} + c_{2,3}) + \alpha c_{3,1} c_{4,1} - \alpha c_{3,2}(c_{2,1} + c_{1,2}) - \beta_3 c_{3,2} - 2\delta c_{3,2} + \delta(c_{1,4} + c_{2,3}) \\
\frac{dc_{1,3}}{dt} &= \beta_3 c_{1,4} + \alpha(c_{1,1} c_{2,2} + c_{1,2} c_{3,1}) - \alpha c_{1,3} c_{4,1} - c_{1,3}(\beta_1 + \beta_2) - 3\delta c_{1,3} + \delta c_{1,4} \\
\frac{dc_{2,3}}{dt} &= \beta_1 c_{1,4} + \alpha(c_{2,1} c_{3,2} + c_{2,2} c_{4,1}) - \alpha c_{2,3} c_{1,1} - c_{2,3}(\beta_2 + \beta_3) - 3\delta c_{2,3} + \delta c_{1,4} \\
\frac{dc_{1,4}}{dt} &= \alpha(c_{1,1} c_{2,3} + c_{1,2} c_{3,2} + c_{1,3} c_{4,1}) - c_{1,4}(\beta_1 + \beta_2 + \beta_3) - 4\delta c_{1,4}
\end{aligned}$$

Assembly of a Heteromeric Chain with Model B Degradation

$$\begin{aligned}
\frac{dc_{1,1}}{dt} &= \beta_1(c_{1,2} + c_{1,3} + c_{1,4}) - \alpha c_{1,1}(c_{2,1} + c_{2,2} + c_{2,3}) - \delta c_{1,1} + Q \\
\frac{dc_{2,1}}{dt} &= \beta_2(c_{2,2} + c_{2,3}) + \beta_1 c_{1,2} - \alpha c_{2,1}(c_{3,1} + c_{3,2} + c_{1,1}) - \delta c_{2,1} + Q \\
\frac{dc_{3,1}}{dt} &= \beta_3 c_{3,2} + \beta_2(c_{1,3} + c_{2,2}) - \alpha c_{3,1}(c_{4,1} + c_{2,1} + c_{1,2}) - \delta c_{3,1} + Q \\
\frac{dc_{4,1}}{dt} &= \beta_3(c_{1,4} + c_{2,3} + c_{3,2}) - \alpha c_{4,1}(c_{3,1} + c_{2,2} + c_{1,3}) - \delta c_{4,1} + Q \\
\frac{dc_{1,2}}{dt} &= \beta_2(c_{1,3} + c_{1,4}) + \alpha c_{1,1} c_{2,1} - \alpha c_{1,2}(c_{3,1} + c_{3,2}) - \beta_1 c_{1,2} - \delta c_{1,2} \\
\frac{dc_{2,2}}{dt} &= \beta_3 c_{2,3} + \beta_1 c_{1,3} + \alpha c_{2,1} c_{3,1} - \alpha c_{2,2}(c_{4,1} + c_{1,1}) - \beta_2 c_{2,2} - \delta c_{2,2} \\
\frac{dc_{3,2}}{dt} &= \beta_2(c_{1,4} + c_{2,3}) + \alpha c_{3,1} c_{4,1} - \alpha c_{3,2}(c_{2,1} + c_{1,2}) - \beta_3 c_{3,2} - \delta c_{3,2} \\
\frac{dc_{1,3}}{dt} &= \beta_3 c_{1,4} + \alpha(c_{1,1} c_{2,2} + c_{1,2} c_{3,1}) - \alpha c_{1,3} c_{4,1} - c_{1,3}(\beta_1 + \beta_2) - \delta c_{1,3} \\
\frac{dc_{2,3}}{dt} &= \beta_1 c_{1,4} + \alpha(c_{2,1} c_{3,2} + c_{2,2} c_{4,1}) - \alpha c_{2,3} c_{1,1} - c_{2,3}(\beta_2 + \beta_3) - \delta c_{2,3} \\
\frac{dc_{1,4}}{dt} &= \alpha(c_{1,1} c_{2,3} + c_{1,2} c_{3,2} + c_{1,3} c_{4,1}) - c_{1,4}(\beta_1 + \beta_2 + \beta_3) - \delta c_{1,4}
\end{aligned}$$

2.6 A note on numerical methods

Numerical integration of the ODEs described in the above sections was carried out in MATLAB 7 [9], using the “ode15s” function with parameters $\text{AbsTol} = 1 \cdot 10^{-15}$ and $\text{RelTol} = 2.22045 \cdot 10^{-14}$. All analytical work was performed using Mathematica 7 [10]. The statistical analyses described in sections 4.1.2 and 5 were carried out using the R statistical computing package [11].

3 Equilibrium and steady-state solutions for the homomeric 3-membered ring

3.1 Equilibrium results without synthesis or degradation

In section 2.1.4, we provided an example of the system of ordinary differential equations representing the dynamics of a simple, homomeric three-membered ring (equation 8). It is straightforward to solve for the concentrations of each of the intermediates at equilibrium. To review, the concentration of the monomer is represented by X_1 , the dimer by X_2 and the trimer by X_3 . We also introduce the variable $X_T \equiv \sum_{i=1}^n i \cdot X_i$ to represent the total concentration of

protein in the system. Setting the three differential equations in equation 8 to 0, we have:

$$\begin{aligned} X_3 &= \frac{X_1 \cdot X_2}{3K_\gamma} \\ X_2 &= \frac{X_1^2}{K_D}, \end{aligned}$$

where $K_\gamma = \gamma/\alpha = \frac{1}{c_0} \cdot K_D^2 \cdot e^{-\Delta G_p^0/RT}$ is used for notational convenience (see section 1). Combining the above relationships with the conservation of mass results in the cubic equation:

$$X_1^3 \left(\frac{1}{K_\gamma \cdot K_D} \right) + X_1^2 \left(\frac{2}{K_D} \right) + X_1 - X_T = 0. \quad (23)$$

Equation 23 can be solved for X_1 trivially using the cubic formula (using Mathematica [10]), with the positive, real solution representing the equilibrium value of X_1 . Given the length of the solution, we do not reproduce the full formula here.

To examine the distribution of mass among the various intermediates, we introduce $Y_i \equiv i \cdot X_i/X_T$ which represents the fraction of total protein in the system found in the intermediate of length i . A plot of the equilibrium values of Y_1 , Y_2 and Y_3 as a function of X_T is shown in Figure 3. There is a transition between monomers and trimers as X_T is decreased; note that the stability of rings is such that this transition occurs at small values of X_T even when the individual affinities are fairly weak ($K_D \sim 10^{-5}$ M). Dimers are always rare at equilibrium ($Y_2 < 10^{-5}$), regardless of X_T .

3.2 Steady-state results including synthesis and degradation

3.2.1 Model A: subunit deletion

In section 2.4.1, we provided an example of model A-type synthesis and degradation for the homomeric three-membered ring (equation 15). To calculate the change in total monomer concentration, X_T , over time, we recall that $X_T = X_1 + 2X_2 + 3X_3$ for a homomeric three-membered ring and find:

$$\begin{aligned} \frac{dX_T}{dt} &= \frac{dX_1}{dt} + 2\frac{dX_2}{dt} + 3\frac{dX_3}{dt} \\ &= Q - \delta X_1 + 2\delta X_2 - 4\delta X_2 + 6\delta X_3 - 9\delta X_3 \\ &= Q - \delta(X_1 + 2X_2 + 3X_3) \\ &= Q - \delta X_T. \end{aligned}$$

This gives the obvious result at steady state:

$$X_T = \frac{Q}{\delta}. \quad (24)$$

To calculate the steady-state concentration of the various species in this case, we begin by




Figure 3: The equilibrium fraction of monomers and trimers for a homomeric three-membered ring as a function of total subunit concentration. “ Y_i ” represents Y_1 or Y_3 as defined in the text. Y_2 is so small at all values of X_T that it cannot be seen here. In this case, we have $K_D = 10^{-5}$, which gives $K_\gamma \sim 10^{-17}$. The transition between monomers and trimers is centered around an X_T of 10^{-10} in this case.

considering X_3 :

$$\begin{aligned} 0 &= \alpha X_1 \cdot X_2 - 3\gamma X_3 - 3\delta X_3 \\ X_3 &= \frac{X_1 \cdot X_2}{3(K_\gamma + K_\delta)}, \end{aligned} \quad (25)$$

where $K_\gamma \equiv \gamma/\alpha$ was defined in section 1 and a new constant $K_\delta \equiv \delta/\alpha$ is introduced. Turning to X_2 , we have:

$$\begin{aligned} 0 &= \alpha X_1^2 - \beta X_2 - \alpha X_1 \cdot X_2 + 3\gamma X_3 - 2\delta X_2 + 3\delta X_3 \\ &= \alpha X_1^2 - \beta X_2 - 2\delta X_2 \\ X_2 &= \frac{X_1^2}{K_D + 2K_\delta}, \end{aligned} \quad (26)$$

where the reduction between the first and second lines is due to the identity in equation 25 above. Combining equation 26 with equation 25 gives:

$$X_3 = \frac{X_1^3}{3(K_\gamma + K_\delta)(K_D + 2K_\delta)}. \quad (27)$$

From the differential equation for X_1 we obtain:

$$0 = 2\beta X_2 - 2\alpha X_1^2 - \alpha X_1 \cdot X_2 + 3\gamma X_3 - \delta X_1 + 2\delta X_2 + Q.$$

Dividing through by δ and recalling equation 24, we obtain:

$$0 = 2 \left(\frac{K_D}{K_\delta} \right) X_2 - 2 \frac{X_1^2}{K_\delta} - \frac{X_1 \cdot X_2}{K_\delta} + 3 \left(\frac{K_\gamma}{K_\delta} \right) X_3 - X_1 + 2X_2 + X_T.$$

Substituting using equations 27 and 26, we obtain the cubic equation:

$$\begin{aligned} 0 = & X_1^3 \left(\frac{K_\gamma}{K_\delta (K_\gamma + K_\delta) (K_D + 2K_\delta)} - \frac{1}{K_\delta (K_D + 2K_\delta)} \right) + \\ & + X_1^2 \left(\frac{2K_D}{K_\delta (K_D + 2K_\delta)} + \frac{2}{K_D + 2K_\delta} - \frac{2}{K_\delta} \right) - X_1 + X_T. \end{aligned} \quad (28)$$

As in section 3.1, it is straightforward to solve the above equation, but due to the length of the formula we do not reproduce it here.

Although equation 28 is true for any set of parameters, recall from section 1 that $\gamma \ll \beta$. For the parameters that we consider here in this work (K_D 's generally lower/stronger than 10^{-5} M and protein half-lives less than 10^8 s), we have $\gamma \ll \delta$. We can obtain a slightly simplified version of equation 28 for the parameters we consider by noting that $K_D + K_\gamma \approx K_D$ and $K_\gamma/K_D \approx 0$:

$$0 = \frac{-X_1^3}{K_\delta (K_D + 2K_\delta)} + X_1^2 \left(\frac{2K_D}{K_\delta (K_D + 2K_\delta)} + \frac{2}{K_D + 2K_\delta} - \frac{2}{K_\delta} \right) - X_1 + X_T. \quad (29)$$

The results displayed in Fig. 2C of the main text, where we plot Y_3 vs. K_D , are calculated using the positive real solution of equation 29.

Given the solutions to either equation 28 or 29, it is natural to ask whether one can obtain a closed-form solution for the value of K_D that maximizes Y_3 by solving the equation $dY_3/dK_D = 0$. Unfortunately, the derivatives in question are extremely complex and as such we have not yet found an analytical solution for the K_D that provides maximum yield.

3.2.2 Model B: whole-complex degradation

A mathematical description of model B-type degradation can be found in section 2.4.3. Here we consider the case of a homomeric three-membered ring, equation 20. As with model A above, we can easily show:

$$\begin{aligned} \frac{dX_T}{dt} &= Q - \delta (X_1 + 2X_2 + 3X_3) \\ &= Q - \delta X_T, \end{aligned}$$

which at steady state yields $X_T = Q/\delta$ as with model A. We first solve for X_3 at steady state:

$$X_3 = \frac{X_1 \cdot X_2}{3K_\gamma + K_\delta}. \quad (30)$$

In the case of model B, degradation of the full ring does not produce a dimer, and so the $\alpha X_1 \cdot X_2$ term in the equation for dX_2/dt does not cancel as was the case for model A. This yields a considerably more complicated equation for X_2 :

$$X_2 = \frac{X_1^2}{X_1 \left(1 - \frac{3K_\gamma}{3K_\gamma + K_\delta}\right) + K_D + K_\delta}. \quad (31)$$

The equation for X_1 at steady state can be written:

$$0 = 2 \left(\frac{K_D}{K_\delta} \right) X_2 - 2 \frac{X_1^2}{K_\delta} - \frac{X_1 \cdot X_2}{K_\delta} + 3 \left(\frac{K_\gamma}{K_\delta} \right) X_3 - X_1 + X_T. \quad (32)$$

Using equations 30 and 31 we can rewrite equation 32, which after simplifying and collecting terms of the same order yields a cubic equation:

$$0 = -3K_\delta X_1^3 - 3K_\delta X_1^2(K_\delta + 2K_\gamma) - K_\delta X_1(K_D(K_\delta + 3K_\gamma) + K_\delta(K_\delta + 3K_\gamma - X_T)) + K_\delta X_T(K_D + K_\delta)(K_\delta + 3K_\gamma). \quad (33)$$

As discussed at the end of section 3.2.1 above, we generally have $K_\gamma \ll K_\delta$. Using this fact, we can obtain a somewhat simpler version of equation 33:

$$0 = -3K_\delta X_1^3 - 3K_\delta^2 X_1^2 - K_\delta^2 X_1(K_D + K_\delta - X_T) + K_\delta^2 X_T(K_D + K_\delta). \quad (34)$$

Analytical results for model B are calculated using the positive real solution of equation 34. As with model A, we have not obtained a closed-form analytical solution for the K_D that maximizes Y_3 using either equation 33 or 34.

Despite their similarities, the dependence of Y_3 on the parameters of the system differs considerably between model A and model B. A detailed analysis of these differences (obtained using the analytical results presented here in addition to results from numerical integration) can be found in section 4.2.2.

3.2.3 Model A degradation for a homomeric chain

As discussed above, for homomeric rings we have found that intermediate affinities can maximize the yield of the full structure at steady state for both of our models of degradation. Here we examine this phenomenon in the case of chain-like structures by considering a simple homodimer under model A degradation. In this case, we have only two types of intermediates (with variables X_1 and X_2 for the concentration of the monomer and dimer, respectively). For this homodimer,

we have the following simple system of ODEs:

$$\begin{aligned}\frac{dX_1}{dt} &= -2\alpha X_1^2 + 2\beta X_2 + Q - \delta X_1 + 2\delta X_2 \\ \frac{dX_2}{dt} &= \alpha X_1^2 - \beta X_2 - 2\delta X_2.\end{aligned}\tag{35}$$

It is simple to verify that, for $X_T \equiv X_1 + 2X_2$, we have $X_T = Q/\delta$ at steady state. From equation 35, we have:

$$X_2 = \frac{X_1^2}{K_D + 2K_\delta}\tag{36}$$

at steady state. Combining this result with the first equation in 35, we get the quadratic equation:

$$0 = X_1^2 \left(\frac{2K_D}{K_\delta(K_D + 2K_\delta)} + \frac{2}{K_D + 2K_\delta} - \frac{2}{K_\delta} \right) - X_1 + X_T$$

for X_1 at steady state. Solving this yields:

$$X_1 = \frac{1}{4} \left(-K_D - 2K_\delta + \sqrt{K_D + 2K_\delta} \cdot \sqrt{8X_T + K_D + 2K_\delta} \right).\tag{37}$$

Unlike the ring models discussed above, it is simple to differentiate Y_2 with respect to K_D to find the K_D at which yield of the homodimer is maximized. We get:

$$\frac{dY_2}{dK_D} = \frac{-4X_T - K_D - K_\delta + \sqrt{K_D + 2K_\delta} \cdot \sqrt{8X_T + K_D + 2K_\delta}}{4X_T \sqrt{K_D + 2K_\delta} \cdot \sqrt{8X_T + K_D + 2K_\delta}},\tag{38}$$

where we have $dY_2/dK_D < 0$ for $X_T \neq 0$. To show this, simply note that $dY_2/dK_D \geq 0$ implies:

$$\begin{aligned}- (4X_T + K_D + 2K_\delta) + \sqrt{K_D + 2K_\delta} \cdot \sqrt{8X_T + K_D + 2K_\delta} &\geq 0 \\ \frac{4X_T}{\sqrt{K_D + 2K_\delta}} + \sqrt{K_D + 2K_\delta} &\leq \sqrt{8X_T + K_D + 2K_\delta} \\ \left(\frac{4X_T}{\sqrt{K_D + 2K_\delta}} \right)^2 + 8X_T + K_D + 2K_\delta &\leq 8X_T + K_D + 2K_\delta \\ X_T^2 &\leq 0.\end{aligned}\tag{39}$$

Since X_T is real, equation 39 can only be satisfied when $X_T = 0$, where dY_2/dK_D is trivially equal to 0 for all values of K_D . Strengthening the interactions (i.e. decreasing K_D) thus *always* increases the yield of the chain. Although this result is very intuitive, it nonetheless provides a counterpoint to the behavior of ring-like structures, where increasing affinity does not necessarily increase the yield of the complex (e.g. Fig. 2C in the main text). Our numerical results on heteromeric chains of length four (section 4.3.3) are consistent with the above results.

4 Additional results

4.1 Ring assembly dynamics and deadlock

4.1.1 Effects of affinity and concentration on ring assembly

As discussed in the main text, homomeric rings with uniformly strong interactions between subunits can exhibit a long-lasting assembly deadlock, or “plateau,” during which formation of the fully assembled ring is blocked due to the persistence of incompatible intermediates (see Fig. 2A in the main text). In this section we further illustrate the effects of affinity and subunit concentration on assembly dynamics.

As shown in Figure 4A, weakening affinities progressively delays the process of ring assembly due to the lack of stability of assembly intermediates. With a monomer concentration of 400 nM, at very weak affinities ($K_D > 10^{-4}$ M) the balance of protein complexes at equilibrium begins to shift away from trimers towards monomers (see also Figure 3 in section 3.1). By contrast, strong interactions induce the plateau, shown in Figure 4B. While increasing affinities progressively increases the duration of the plateau phase, the plateau always halts assembly at the same concentration of fully assembled complex—that is, the plateau “height” is invariant for a particular ring size. These results demonstrate that for ring assembly efficiency, there exists an *optimal* affinity for any given ring size and subunit concentration.




Figure 4: Effect of affinity on ring assembly dynamics. Shown are the results of numerical simulation of the ordinary differential equations for the three-membered homomeric ring (equation 8). $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$, initial monomer concentration of 400 nM. (A) Weaker-than-optimal interactions. (B) Stronger-than-optimal interactions.

The initial concentration of monomers also affects the ring assembly dynamics, as shown in Figure 5. Reducing the monomer concentration delays assembly, due to the lower frequency of interaction between subunits (e.g., the red curve). Increasing the monomer concentration induces a plateau phase, but, importantly, higher concentrations induce the plateau at progressively *earlier* timescales. While this allows for shorter waiting times to reach percentage yields smaller than the plateau percentage (purple curve), it ultimately does not affect the time taken to achieve yields greater than the plateau (overlap of the purple, blue and green curves). By contrast,

increasing affinities while keeping monomer concentration constant induces a plateau phase that progressively *delays* assembly for these higher yields, while leaving the time taken to reach the sub-plateau yields completely unaffected (Figure 4B).




Figure 5: Effect of initial monomer concentration on ring assembly dynamics. Numerical simulations of the ordinary differential equations for the three-membered homomeric ring (equation 8). $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$, $\beta = 2.53 \cdot 10^{-3} \text{ s}^{-1}$, $K_D = \beta/\alpha = 10^{-9} \text{ M}$.

The different role played by affinity vs. concentration in ring assembly dynamics can be explained as follows. Since the plateau phase results from the persistence of incompatible intermediates, the time of *initiation* of the plateau phase is determined by the rate of complex formation, which is in turn determined by the on rate α and the concentration of monomers. However, the *resolution* of the plateau phase is determined by the timescale of dissociation of the interactions: the more stable the incompatible intermediates are, the longer the average waiting time until they dissociate and return subunits to the pool.

These observations also allow us to make the prediction that in experimental assays of the assembly of homomeric rings (e.g., of the type described by Kress et al. [12]), one should always be able to induce the plateau by increasing the monomer concentration. Actually observing the plateau, however, requires that the time resolution of the experiment be sufficiently small, which may prove difficult in certain cases. It is also important to note that the plateau effect will be difficult to observe unless the data from the assay is plotted on a logarithmic time scale.

4.1.2 Scaling of plateau height with ring length

As mentioned in the main text, we find that the length of the ring in question has a strong influence on the concentration of fully formed rings observed during the “plateau” phase of the dynamics. We find that the value of $Y_n = n \cdot X_n/X_T$ observed at the plateau is an invariant function of concentration for fixed ring length (Figures 4 and 5) and that increasing ring length

decreases Y_n in the plateau phase (Figure 6).




Figure 6: Effect of ring length on ring assembly dynamics. Numerical simulations of the ordinary differential equations for homomeric ring assembly described in section 2.1.3. $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$, $K_D = 10^{-12} \text{ M}$, initial monomer concentration of 400 nM. The plateau occurs at progressively lower yields as ring length increases.




Figure 7: Effect of ring length on plateau height. Both panels represent data from numerical simulations of the ordinary differential equations for homomeric ring assembly described in section 2.1.3. $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$. Initial monomer concentrations set to $n \cdot 40\text{nM}$. Plateaus calculated numerically by sampling the fraction of full complex at $t = 10^8$ seconds in the presence of very strong interactions ($K_D = 10^{-24} \text{ M}$). (A) Here we plot Y_n directly for rings of length 3 to 22. Note the difference between rings with odd and even lengths. (B) Here we plot the same data as in (A), but using the inverse of the plateau height (Y_n^{-1}) for the y coordinate. The solid lines represent the best linear fits to the data, with $Y_n^{-1} = 0.320n + 0.521$ for odd rings and $Y_n^{-1} = 0.320n + 0.430$ for even rings.

The regular relationship between “plateau height” and ring length is shown in Figure 7.

Interestingly, we observe that odd and even rings have slightly different scaling (Figure 7A). In both cases, the relationship between Y_n^{-1} and n is approximately linear at the plateau (Figure 7B). For odd rings, we have $Y_n^{-1} = 0.320n + 0.521$ with $p < 2 \cdot 10^{-16}$ for both coefficients. For even rings, we have $Y_n^{-1} = 0.320n + 0.430$, again with $p < 2 \cdot 10^{-16}$ for both coefficients. Although the difference in intercepts is slight (Figure 7B), we find that the intercept terms are likely distinct, in that their 95% confidence intervals do not overlap. The slope terms, however, are statistically indistinguishable.

At present, we have not determined an exact analytical relationship that can explain the scaling of plateau height with n that we observe.

4.1.3 Assembly time as a function of affinity, ring length, and concentration

The dependence of ring assembly dynamics on affinity, concentration, and ring length, described in sections 4.1.1 and 4.1.2, can be summarized by plotting assembly time as a function of these three variables, as shown in Figure 8. As described in the main text we describe assembly efficiency using T_X . This is the amount of time it takes a system that starts out with 100% monomers to reach $X\%$ of the maximum concentration of the full ring. Each panel within Figure 8 consists of a series of curves showing the assembly time of the ring as a function of affinity; each curve within the panel represents a starting monomer concentration, as denoted in the figure legend. The first row of panels represent the results for a ring of length three; the second row, a ring of length seven. The columns of the plot show the assembly-time curves for target yields of 99%, 75% and 50% (i.e., T_{99} , T_{75} , and T_{50}).

The plots in Figure 8 highlight a number of the dynamic effects described in depth in sections 4.1.1 and 4.1.2 above:

- In each plot, the curves have distinct minima at which assembly time is optimized for that ring length and concentration. Assembly time increases when affinities are either stronger or weaker than this optimal level. The notable exception to this pattern is the T_{50} curve for the three-ring, in which each curve is nearly flat when interactions are strong. This is due to the fact that for the three-ring, 50% yield is reached *before* the plateau phase begins, so affinity has minimal effect on the T_{50} (see Figure 4B). By contrast, for the ring of length seven, T_{50} increases when interactions are strong: this is due to the fact that the plateau phase has already occurred (at approximately 36% yield—see Figures 6 and 7A), and hence its duration becomes rate-limiting for assembly measured by T_{50} .
- Generally, increasing concentration decreases assembly time by increasing the rate of association between subunits (violet curves are always lower than red curves). However, when interactions are very strong, increasing concentration has minimal effect on assembly time because the rate of dissociation of the incompatible intermediates during the plateau phase becomes the rate-limiting process (see Figure 5).
- For larger rings, having weaker-than-optimal interactions incurs a greater “penalty” in terms of assembly time than for shorter rings. This can be seen in the much greater upward slope of the families of curves at weaker-than-optimal interactions for rings of length seven.

This phenomenon is due to the fact that, when weak affinities are the limiting factor in assembly, the larger intermediates that must form during the assembly of larger rings are more likely to break apart before finding a suitable binding partner to form the full ring.




Figure 8: Assembly time as a function of affinity, concentration, and ring length. Times to reach 99%, 75% or 50% yield (T_{99} , T_{75} , and T_{50} , respectively) were calculated by numerical simulation of the ordinary differential equations for homomeric ring assembly described in section 2.1.3. $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$ for all simulations.

An additional aspect of the dynamics that can be seen in Figure 8 is that the affinity yielding optimal concentration (the minimum of each curve) changes as a function of the concentration. This is shown clearly for the homomeric three- and seven-rings in Figure 9. The relationship between the K_D yielding the fastest assembly (lowest T_{99}) and the initial monomer concentration is linear with a numerically calculated relationship of $y = \frac{1}{40}x$ for the three-ring, and $y = 4.457 \cdot 10^{-3}x - 4.521 \cdot 10^{-14}$ for the seven-ring.

4.2 Synthesis and degradation

4.2.1 Ring assembly dynamics with synthesis and degradation

As discussed in sections 2.4 and 3.2, we also explored the effect of synthesis and degradation on the dynamics of ring assembly. In this section we include additional numerical and analytical results from these models.

Figure 10 shows the assembly dynamics of the homomeric three-ring in the case of model A




Figure 9: Optimal affinity as a function of monomer concentration. For each concentration shown, results were calculated by varying the K_D in increments of $10^{\frac{1}{4}}$, performing numerical simulation of the differential equations given in section 2.1.3, and identifying the K_D leading to the lowest T_{99} . $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$. The plot is scaled logarithmically to show the entire range of numerical data. A fit yields a slope of 1, that is, a linear relationship between optimal K_D and monomer concentration, as detailed in the text.

(Figure 10A) and model B (Figure 10B) degradation. In the presence of synthesis and degradation, assembly reaches a steady state where the yield of the full ring is less than 100%, and the concentrations of intermediate species are non-zero (compare to Fig. 2A in the main text). Furthermore, the steady-state yield of the ring is higher in the presence of model B degradation than for model A degradation, using nominally identical rates for synthesis and degradation.




Figure 10: Ring assembly in the presence of synthesis and degradation. Results based on numerical simulation of the ordinary differential equations for the homomeric three-ring described in sections 2.4.1 and 2.4.3. The parameter values for both the simulation in both plots are identical. As described in the main text, the parameter values $X_T = Q/\delta = 477 \text{ nM}$ and $\delta = 2.75 \cdot 10^{-4} \text{ s}^{-1}$ are taken from averages in available experimental data on yeast proteins [13, 14]. From these numbers, we can calculate $Q = X_T \cdot \delta = 1.31 \cdot 10^{-10} \text{ M s}^{-1}$. For these simulations, we also have $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$, $K_D = 10^{-12} \text{ M}$. (A) Degradation model A. (B) Degradation model B.

4.2.2 Effect of K_D and synthesis rate on assembly yield

As with assembly time (see Figure 4), the steady-state yield of the full complex in the presence of synthesis and degradation is sensitive to the affinities between subunits. Figure 11 shows the effect of varying the affinity on the dynamics of formation of the full ring, in the presence of model A-type degradation. Going from very strong interactions to intermediate interactions, the steady-state yield improves (red vs. green curve). However, further weakening of the interactions rapidly diminishes yield (green vs. blue curve). This result indicates that in the presence of synthesis and degradation, interactions that are weaker or stronger than an optimal value will diminish the yield of the ring.




Figure 11: Assembly dynamics of full ring vs. affinity in the presence of synthesis and degradation. Degradation model A, other parameters as for Figure 10.

In the case of the homomeric three-ring, we can explore this effect both numerically and analytically (see section 3.2 and Fig. 2C in the main text). The results of this analysis for both model A and model B are shown in Figure 12. As also demonstrated in Figure 11, very strong interactions can suppress the assembly of the full complex, though the extent of this effect is dependent on the model for degradation that is assumed (see also Figure 14).

These differences between the degradation models can be understood by referring to the timecourses of assembly in the presence of strong interactions, shown in Figure 10. In model A degradation (Figure 10A), the likelihood of a complex experiencing a deletion event increases in proportion with its length, since the deletion rate δ applies to subunits rather than complexes. The fully assembled rings are thus subject to a higher rate of degradation than lower-order intermediates. In addition, deletion of a subunit from a trimer produces a dimer; since dimers cannot interact to form trimers and will tend to persist in the presence of strong interactions, model A degradation leads to a steady state in which dimers occupy a relatively larger fraction of total mass, at the expense of trimers. In model B degradation (Figure 10B), monomers, dimers, and trimers are all equally likely to experience degradation, and degradation of a trimer removes it completely from the system rather than giving rise to an assembly-incompatible dimer. Therefore when the model B system has reached steady state it produces a relatively larger yield of trimers than model A, and a lower balance of dimers. Since these simulations were performed




Figure 12: Effect of binding strength on steady state yield of fully assembled complex in the presence of synthesis and degradation. As in Figure 10, $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$, $Q = 1.31 \cdot 10^{-10} \text{ M s}^{-1}$, $\delta = 2.75 \cdot 10^{-4} \text{ s}^{-1}$, $X_T = Q/\delta = 477 \text{ nM}$. In both cases, we can consider these curves to represent the effect of varying affinity on rings formed from “average” yeast proteins [13, 14]. (A) Model A degradation. Solid line indicates analytical solution of equation 29; circles represent results from numerical integration of equation 15. (B) Model B degradation. Solid line indicates analytical solution of equation 34; circles represent results from numerical integration of equation 20. Note difference in scale between (A) and (B).

with strong interactions, the steady state concentration of monomers is fairly low in both cases.

The steady state yield of the full ring depends not only on the affinity, but also on the rates of synthesis and degradation as well as the total concentration of protein in the system. In an analogous fashion to Figure 12, Figure 13 shows plots of yield vs. K_D , but here each curve represents a different value for the total amount of mass in the system X_T , produced by changing the synthesis rate Q while holding the degradation rate δ constant (changing Q and not δ allows us to look specifically at the effect of changes in X_T while leaving the parameter K_δ , the ratio between the degradation rate δ and the association rate α , unchanged; see section 3.2). For each curve, the yield Y_n is calculated relative to the total amount of protein at steady state, given by Q/δ (see section 3.2).

The curves in Figure 13 become higher for increasing X_T , indicating that for any given affinity, increasing the total amount of protein also increases the proportion of mass that assembles into the full ring. In addition, the effect of affinity optimality also changes as a function of X_T . At very *low* X_T , degradation, rather than deadlock, is the limiting factor in assembly, and hence alleviating deadlock via affinity optimization has minimal effect in boosting steady state yield. At very *high* X_T , the larger influx of new monomers quickly alleviates deadlock, which again mitigates the effect of overly strong interactions (purple curve). At *intermediate* synthesis rates, strong interactions lead to a significant fraction of deadlocked intermediates at steady state, which is alleviated by weakening interactions until an optimal yield is reached (green curve). However, beyond the point of optimality, weakening affinities leads to reduced yield, regardless of the total amount of protein in the system.




Figure 13: Steady state yield as a function of K_D and total protein concentration, model A degradation. Results generated by numerical simulation of the ordinary differential equations described in section 2.4.1; different values of X_T were obtained by changing the value of the synthesis rate Q while holding δ at a constant value of $2.75 \cdot 10^{-4} \text{ s}^{-1}$. $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$.

The magnitude of the affinity optimality effect for both degradation models is summarized in Figure 14. This figure tracks the increase in yield that can be achieved by using optimal interactions rather than very strong interactions, for a range of values of X_T . As discussed above for Figure 13, the improvement in yield that can be gained by using optimal interactions is dependent on the total amount of protein X_T , and reaches a maximum of approximately 7.5% for model A at $X_T = 24 \text{ nM}$ with $\delta = 2.75 \cdot 10^{-4}$. In addition, as discussed for Figure 12, model B degradation exhibits a much less pronounced affinity optimality effect for all values for X_T , with a maximal improvement in yield of only 0.51% at $X_T = 6.0 \text{ nM}$.




Figure 14: Improvement in assembly yield at optimal K_D vs. very strong interactions ($K_D = 10^{-13}$) at different values of X_T . Results generated by numerical simulation of the ordinary differential equations described in sections 2.4.1 and 2.4.3; different values of X_T were obtained by changing the value of the synthesis rate Q while holding δ at a constant value of $2.75 \cdot 10^{-4} \text{ s}^{-1}$. $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$.

Though the results shown in Figures 13 and 14 help to elucidate the underlying dynamics, our analytical solutions for the homomeric three-ring for both model A (equation 29) and model B (equation 34) reveal that the steady state yield is a complex function not only of K_D and X_T , but also of the ratio of the degradation rate to the association rate, K_δ . We leave a thorough characterization of the relative contributions of these different parameters to future work. However, taken together, these results indicate that ring assembly efficiency can be inhibited by uniformly strong interactions in the presence of synthesis and degradation, though the magnitude of this effect depends on the particular model of degradation that is assumed and the choice of synthesis and degradation parameters.

4.3 Effect of affinity configurations on assembly

To determine the effect of affinity configuration on the assembly efficiency of heteromeric rings, we systematically sampled the space of possible configurations and performed numerical simulations of assembly for each one. In this section we describe our methods for enumerating configurations and present additional results for the assembly of 4-, 5-, and 6-membered heteromeric rings.

4.3.1 Enumerating distinct affinity configurations for heteromeric rings

To consider a large space of possible affinity configurations, we allowed the affinity at each subunit interface to vary over seven orders of magnitude, with a K_D of either 10^{-12} , 10^{-11} , 10^{-10} , 10^{-9} , 10^{-8} , 10^{-7} , or 10^{-6} M. If each subunit interface were considered to be distinct, this would imply $7^3 = 343$ unique configurations for the homomeric three-ring, $7^4 = 2401$ for the four-ring, and so on. However, as shown in Figure 15, such an approach would include many redundant configurations differing only in the (arbitrary) labeling of the subunits. For example, having a *strong* interaction between x_0 and x_1 , a *medium* strength interaction between x_1 and x_2 , and a *weak* interaction between x_2 and x_0 (a configuration we will denote $S - M - W$, with the abbreviated affinities enumerated in subunit counting order—see section 2.2.1), is equivalent in its assembly properties to the configuration having a *medium* strength interaction between x_0 and x_1 , a *weak* interaction between x_1 and x_2 , and a *strong* interaction between x_2 and x_0 (a configuration we denote $M - W - S$, using the same convention). The equivalence arises from the fact that the $S - M - W$ arrangement of interactions can be converted into the $M - W - S$ arrangement simply by shifting the labels of the subunits around the ring one position while preserving their sequential ordering. By thus considering the *rotational* symmetry of the interactions around the ring we can ignore these dynamically identical configurations.

In addition, since the affinity configurations under consideration are two-dimensional and do not distinguish a “top” or “bottom” for the ring, to identify unique configurations we must also note that the subunit counting direction—clockwise vs. counter-clockwise—also does not affect the dynamics. Changing the subunit counting direction is equivalent to *reflecting* the subunit labels about an axis in the plane, which can also be imagined as “flipping” the ring structure over. As shown in Figure 15, if the subunit labels for the ring with the affinity configuration $S - M - W$ are “flipped over” or reflected, it results in the nominally distinct configuration $W - M - S$; while the counting order of the interactions has changed from clockwise to counter-clockwise, the




Figure 15: Ring symmetries. A unique affinity configuration can be given several different names depending on how subunits are labeled or counted. To identify all unique affinity configurations, all possible notationally distinct configurations are considered and then redundant configurations are eliminated. Redundant configurations are identified by their equivalence to an existing configuration via rotational symmetry, corresponding to a change in the placement of the first subunit label x_0 along the ring, or by symmetry of reflection, corresponding to a change in the direction in which subunits are enumerated. By considering these two types of symmetry, the six notationally distinct affinity configurations shown can be seen to be equivalent, and represented by a single member, $S - M - W$.

assembly dynamics of the configuration have not changed. Thus, nominally distinct configurations that are identical by the symmetries of rotation and reflection can be ignored, and only one representative of the family of equivalent configurations need be considered for analysis.

Proceeding in this fashion, we enumerated the unique affinity configurations for the 3-, 4-, 5-, and 6-membered rings, allowing the K_D at each interface to vary over seven orders of magnitude as described above. With the exclusion of configurations that are redundant by symmetry, we obtain a significantly reduced number of possible configurations (see Table 1).

Ring Length	Unique Configurations
3	81
4	406
5	1,855
6	10,528
7	60,028

Table 1: Numbers of distinct affinity configurations for rings of different lengths after accounting for symmetries (see Figure 15).

4.3.2 Numerical simulation results for the heteromeric 3-, 4-, 5-, and 6-rings

For each of the affinity configurations enumerated as above we ran numerical simulations of heteromeric ring assembly using the systems of ordinary differential equations described in sections 2.2.2, 2.4.2 and 2.4.4.

For the three-ring (Fig. 3A in the main text, reproduced here as Figure 16), the results show that configurations involving either one or two weak interactions assemble more efficiently than configurations involving uniform interactions. As shown in the inset plots, the differences in efficiency between configurations with one or two weak interactions are very small in absolute terms. It is also worth noting that the relative T_{99} rankings of one- versus two-weak interaction configurations is dependent in part on monomer concentration, while the model A and model B rankings are not.




Figure 16: Ranking the assembly efficiency and yield of affinity configurations for the heteromeric three-ring. This figure corresponds to Fig. 3A in the main text and is included here for completeness. For calculation of T_{99} , simulations were performed with initial monomer concentrations of 477 nM for each subunit. For simulations of model A and model B degradation, $Q = 1.31 \cdot 10^{-10} \text{ M s}^{-1}$ and $\delta = 2.75 \cdot 10^{-4} \text{ s}^{-1}$. $X_T = Q/\delta = 477 \text{ nM}$. For all simulations, $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$. “One Weak” denotes a configuration with binding strengths (i.e., K_D s) of 10^{-12} , 10^{-12} , and 10^{-6} M . “Two Weak”: 10^{-12} , 10^{-6} , and 10^{-6} M ; “All Medium”: 10^{-8} , 10^{-8} , and 10^{-8} M ; “All Weak”: 10^{-6} , 10^{-6} , and 10^{-6} M ; “All Strong”: 10^{-12} , 10^{-12} , and 10^{-12} M .

For the four-ring (Figure 17), the configuration containing alternating strong and weak interactions outranked other configurations based on T_{99} and model A yield; for model B, however, the configuration with a single weak interaction resulted in the highest yield. Configurations with either one or two weak interactions (with the rest strong) outperformed configurations with uniform interactions or more than two weak interactions.

The results for the five-ring (Figure 18) were similar to those for the four-ring: again, configurations with two weak interactions assembled with the lowest values for T_{99} , slightly ahead of the one-weak interaction configuration. However, the one-weak interaction configuration




Figure 17: Ranking the assembly efficiency and yield of affinity configurations for the heteromeric four-ring. Parameters as for Figure 16. “S”, “M”, and “W” denote strong ($K_D = 10^{-12}$ M), medium ($K_D = 10^{-8}$ M), and weak ($K_D = 10^{-6}$ M) interactions, respectively.

produced the highest model B yield. Configurations with uniform affinities assembled both with low efficiency and low yield. Interestingly, changing the strength of one interaction in the configuration $S - S - S - W - W$ from strong to medium (resulting in the configuration $S - S - M - W - W$) dramatically reduces its performance relative to the other configurations.




Figure 18: Ranking the assembly efficiency and yield of affinity configurations for the heteromeric five-ring. Parameters as for Figure 16; “S”, “M”, and “W” denote strong ($K_D = 10^{-12}$ M), medium ($K_D = 10^{-8}$ M), and weak ($K_D = 10^{-6}$ M) interactions, respectively.

The results for the six-ring (Figure 19) also indicate the assembly efficiency of rings with either one or two weak interactions. Both configurations with two weak interactions slightly outperform the one-weak interaction configuration for T_{99} and model A yield, while the one-weak interaction produces a higher model B yield. All three outperform uniform configurations or configurations with more than two weak interactions. Interestingly, the configuration with alternating strong and weak interactions does not perform as well by any metric as the configurations with either one or two weak interactions. This result, along with the results for the four-ring, in which alternating strong-weak interactions performed very well, suggests that the key to assembly efficiency is not alternating strong and weak interactions in and of itself, but the presence of *two* weak interactions. This may be due to the fact that a ring with two weak interactions may be able to assemble in two halves that then “snap” together to form a stable ring.




Figure 19: Ranking the assembly efficiency and yield of affinity configurations for the heteromeric six-ring. Parameters as for Figure 16. “S”, “M”, and “W” denote strong ($K_D = 10^{-12}$ M), medium ($K_D = 10^{-8}$ M), and weak ($K_D = 10^{-6}$ M) interactions, respectively.

4.3.3 Optimizing the assembly of chains

To provide a contrast to our results for heteromeric three-membered rings (Figure 16), we considered how affinity configuration influences the assembly efficiency of four-membered heteromeric chains. Although these two structures share the same number of interactions, chains lack the rotational symmetry of rings, and as such there are over twice as many distinct

configurations in this case (196 for the chains vs. 81 for the rings). Given some affinity configuration (and its attendant dissociation rates β_1 to β_3), we performed numerical simulations of simple heteromeric chain assembly, assembly with model A degradation, and assembly with model B assembly (the equations used for these calculations are listed in section 2.5).

The relative performance of these affinity configurations is compared in Figure 20. Since chains are generally much less stable than rings (see section 1), most of the configurations sampled here did not assemble to a level of 99%, making it impossible to define T_{99} as for rings. In simulations that did not consider synthesis and degradation, we thus calculated equilibrium yield instead (“Eq. Yield” in Figure 20). Note that, in stark contrast to the behavior for rings, for chains stronger is always better, regardless of the efficiency metric in question.




Figure 20: Ranking the assembly efficiency and yield of affinity configurations for the heteromeric 4-membered chain. Given that most affinity configurations do not assemble to 99% at the concentration considered here, in this case we use equilibrium yield (the ‘Eq. Yield’ column to the far left of the ranking plot) rather than T_{99} to characterize assembly efficiency in the absence of synthesis and degradation. The equilibrium yields were calculated at a total concentration of 477 nM for each monomer type. For simulations of model A and model B degradation, $Q = 1.31 \cdot 10^{-10} \text{ M s}^{-1}$ and $\delta = 2.75 \cdot 10^{-4} \text{ s}^{-1}$. $X_T = Q/\delta = 477 \text{ nM}$. For all simulations, $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$. “One Weak” denotes a configuration with binding strengths (i.e., K_{DS}) of 10^{-12} , 10^{-12} , and 10^{-6} M . “Two Weak”: 10^{-12} , 10^{-6} , and 10^{-6} M ; “All Medium”: 10^{-8} , 10^{-8} , and 10^{-8} M ; “All Weak”: 10^{-6} , 10^{-6} , and 10^{-6} M ; “All Strong”: 10^{-12} , 10^{-12} , and 10^{-12} M .

4.4 Hierarchical assembly pathways

As discussed in section 4.3 above, introducing weak interactions into heteromeric ring structures can dramatically improve their assembly efficiency according to a wide variety of measures. In this section we compare an alternative mechanism by which kinetic assembly bottlenecks can be addressed: namely, the sequential, stepwise assembly of subunits to form the ring. In the case of the heteromeric three-ring, this type of assembly would imply that, for example, x_0 can bind x_1 ,

but the binding of x_1 to x_2 is contingent on x_1 's being bound to x_0 , and the binding of x_2 to x_0 is similarly contingent on x_2 's being bound to x_1 . This model represents a case where allosteric interactions—or other sources of hierarchical structure—force assembly to proceed sequentially, thus preventing the accumulation of assembly intermediates that are incompatible. Yin and coworkers have deployed precisely this approach to optimize the assembly process of cyclic DNA nanostructures [15]. To compare the assembly efficiency of this allosteric approach with the biophysical strategies described in section 4.3, we created and analyzed a model of stepwise assembly as described below.

4.4.1 Mathematical model of stepwise assembly

Before proceeding, we describe here the set of ordinary differential equations used to model the assembly of a heteromeric three-ring via sequential, stepwise assembly. We use the notational convention for heteromeric rings described in section 2.2.1. In this highly simplified model, there is no binding between x_1 and x_2 , or between x_2 and x_0 , unless x_0 has first bound to x_1 . As a result, the concentrations of the dimers $x_{1,2}$ and $x_{2,2}$ are zero for all time. In addition, because the ring breakage rate is much smaller than the dissociation rate (i.e., $\gamma \ll \beta$) for the parameter values we consider below, for simplicity we ignore the process of ring breakage in this case (i.e. we set $\gamma = 0$).

The ODEs are as follows:

$$\begin{aligned}
\frac{dx_{0,1}}{dt} &= \beta x_{0,2} - \alpha x_{0,1} x_{1,1} \\
\frac{dx_{1,1}}{dt} &= \beta x_{0,2} - \alpha x_{0,1} x_{1,1} \\
\frac{dx_{2,1}}{dt} &= -\alpha x_{2,1} x_{0,2} \\
\frac{dx_{0,2}}{dt} &= \alpha x_{0,1} x_{1,1} - \beta x_{0,2} - \alpha x_{0,2} x_{2,1} \\
\frac{dx_3}{dt} &= \alpha x_{2,1} x_{0,2}.
\end{aligned} \tag{40}$$

4.4.2 Comparing sequential assembly with weak interactions

To compare between sequential assembly and the biophysical strategies discussed in section 4.3, we numerically integrated the ODEs from equation 40. In Figure 21 we compare the assembly dynamics of this sequential model, a ring containing a single very weak interaction, and a ring with uniformly strong interactions. In this case we have chosen affinity configurations such that the uniform case and the single weak interaction case exhibit identical thermodynamic stabilities. We find that the single weak interaction configuration always assembles faster than the sequential case, although the magnitude of this difference varies with total monomer concentration.




Figure 21: Assembly timecourses for three-membered heteromeric rings with uniform interactions (K_{DS} of 10^{-9}), one weak interaction (K_{DS} of 10^{-12} , 10^{-12} , and 10^{-3}), and sequential assembly (uniform K_{DS} of 10^{-9}). $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$. For simplicity, the ring breakage rate γ is set to zero in all cases. (A) Monomer concentration of $40 \mu\text{M}$. (B) Monomer concentration of 0.4 nM . Although the three cases here have equivalent thermodynamic stability, they display different assembly kinetics. Note that the “one weak interaction” configuration consistently assembles faster than the sequential case, though the magnitude of this difference depends on concentration.

The differences observed in Figure 21 can be best understood in terms of a schematic view of the assembly process, Figure 22. In this illustration, we consider a pool of monomers A , B , and C that interact with each other to form a heteromeric ring. For simplicity, we do not explicitly consider trimer formation in this schematic, focusing entirely on the process of monomers binding to form dimers.

For the non-sequential assembly models (Figure 22, panels A and B), there are three “paths” that the system can take to assemble full rings from a starting pool of monomers: 1) bind A and B , then C ; 2) bind A and C , then B ; 3) bind B and C , then A . Configurations with uniformly strong interactions attempt to take all three paths simultaneously (represented by the fact that all possible dimers are present in Figure 22A). When concentrations are high, the system consumes all possible monomers too quickly, and since the interactions are strong, a plateau is induced (as discussed above and in the main text). During the plateau phase, assembly via any given path can only proceed when the system “backtracks” from one of the other paths through dissociation of a dimer. Uniformly strong interactions thus lead to unavoidable deadlocks at high concentrations (Figure 21A).

When one of the interactions along the ring is weak, although the system can in theory take all three possible assembly paths, only two of those paths will actually be taken by the majority of proteins in the system. If we make the $A - C$ interaction weak, any monomers attempting to take path number “2” by first forming an $A - C$ interaction will ultimately be unsuccessful, since this interaction will tend to break soon after formation. This fact is represented schematically in Figure 22B by the lack of $A - C$ dimers. As the schematic demonstrates, any monomers that take a given path (say, by forming the stable $A - B$ dimer) are guaranteed to have access to the




Figure 22: Schematic diagram of ring assembly. Shown are the interactions that would form after an arbitrary amount of time given a pool of monomers at equal concentrations and equivalent association rates. For simplicity, only the formation of dimers is considered here. (A) Uniform interactions: all possible dimers can form. (B) One weak interaction: $A - B$ and $B - C$ dimers can form, but encounters between A and C do not result in a stable interaction. (C) Sequential assembly: only encounters between A and B result in a stable interaction. Since the association rates are the same as in (A) and (B), the number of $A - B$ dimers formed is the same (two of the four possible).

cognate monomer needed to complete assembly (C in this example). Configurations with a single weak interaction thus avoid the problem of deadlock at high concentrations and achieve efficient assembly across a wide variety of conditions (see Fig. 3 in the main text).

In the sequential case, assembly can only proceed down a single path—in our example, this is path “1” (Figure 22C). This strategy avoids the potential for deadlock, and as such we do not see a plateau here (Figure 21A). The fact that sequential assembly occurs more slowly than “single weak interaction” assembly arises from the fact that the latter can take two paths concurrently, while the former is forced to take only one. Since the association rate is assumed to be identical for both cases, weak interaction assembly initially exhibits twice the number of productive reactions (a fact schematized by the existence of both $A - B$ and $B - C$ dimers in Figure 22B vs. only $A - B$ dimers in Figure 22C).

We thus find that inclusion of a single weak interaction in a “concurrently” assembling trimer provides the best of both worlds—fast assembly at low concentrations while avoiding deadlock at high concentrations. It is important to note that our analysis here is not exhaustive, and there may be conditions in which sequential assembly will be favored over the concurrent case. We leave further exploration of these alternative approaches to future work.

4.5 Non-uniform concentrations

Since each protein in a heteromeric ring will be transcribed and translated from a separate gene, such rings have the capacity to not only demonstrate variation in affinities (as discussed above) but also the total concentration of each subunit. Such differences could easily arise from the inherent stochastic nature of gene expression [16], or from differences in gene regulation between subunits in the cell. To test the potential effects of such variation, we considered the assembly dynamics of a three-membered heteromeric ring in a situation where one subunit has a higher total concentration than the other two.




Figure 23: Assembly dynamics of a three-member heteromeric ring with non-uniform subunit concentrations. The black curve represents a case where all subunits are in stoichiometric concentrations (i.e. the situation considered extensively above and in the main text). The red curve represents a case where a single subunit is at double the concentration of the other two, and the green curve represents a case where one subunit is triple the concentration of the other two. As one can see, as the difference in concentrations increases, the plateau becomes “lower” and persists longer (i.e. T_{99} increases dramatically). In all plots, two of the subunits are present at $4 \mu\text{M}$ concentration, while the other varies from 4 to $12 \mu\text{M}$. The affinities in this case are uniform with K_{DS} of 10^{-9} M , and the parameter $\alpha = 2.53 \cdot 10^6 \text{ M}^{-1} \text{ s}^{-1}$.

As one can see from Figure 23, increasing the concentration of a single subunit exacerbates deadlock, resulting in a deadlocked plateau that occurs at a lower assembly yield and that persists longer. This occurs because the subunit that is at higher concentration (say, the “A” subunit of a “ABC” heteromeric ring) rapidly binds to the other two subunits, forming a comparatively large number of AB and AC dimers that must dissociate in order for assembly to proceed to completion. These findings highlight the fact that the dynamics of assembly in this case depend not solely on K_{DS} (i.e. the free energy of binding) but also on subunit concentrations (which influence the *chemical potential* of the bimolecular reactions in question). It is currently unclear if either hierarchical assembly or simple affinity configurations can overcome the increased deadlock

resulting from non-uniform concentrations. We leave complete consideration of how chemical potential landscapes might evolve to manage differences in subunit concentration to future work.

5 Analysis of structural data

5.1 Structures for heteromeric three-membered rings

Our work strongly implies that structures that include at least one weak interaction will enjoy an evolutionary advantage, either in terms of assembly time or steady-state yield (Fig. 3A of the main text). We explored the solved structures of ring-like complexes in order to assess if there was any evidence for the existence of such weak interactions in the ring-like complexes found in living systems.

We began by assembling a set of three-membered heteromeric rings of known structure. We used the database 3D Complex [17] as a starting point for collecting the structures, and based our analysis on all heteromeric three-membered complexes in that database with a ring-like topology. In order to reduce the redundancy of the data set (i.e. to avoid considering two very closely related or identical rings as different examples of evolutionarily optimized structures), we utilized the “QS-90” level of the 3D Complex hierarchy [17]. At QS-90, complexes with greater than 90% sequence identity are grouped together into a single class, from which a single representative structure is taken. Using the QS-90 level of the hierarchy allows us to ignore cases where multiple mutant forms of the same complex, or very closely related complexes, have solved structures in the PDB.

We curated the resulting 82 heteromeric three-rings in the database in order to remove structures in which the biology of assembly did not match the case considered by our model. Specifically, we removed structures in the following four classes:

1. **Antibody-Antigen Complexes** Antibodies consist of two polypeptide chains (Heavy and Light) that interact extensively with each other. In many cases, both chains interact with an antigen, thus forming a ring-like topology. Biologically, however, antibodies are synthesized and secreted in the absence of antigen, and only then bind to the antigen in question. Our model does not cover this case, and so we do not consider this type of complex in our data set. It is important to note, however, that the interactions to the antigen generally involve much smaller surface area than the Heavy-Light interaction, but even though these structures support the conclusions of our model we cannot be certain that the evolutionary pressures on this system are equivalent to those implied by our model. Of the 82 ring structures in the initial dataset, 28 belonged to this class.
2. **Integral Membrane Complexes** The vast majority of these cases involve the extracellular domains of dimeric membrane-bound receptors binding to monomeric cytokines. This situation, in which two members of the complex are constrained to a membrane surface and one can diffuse in three-dimensional space, presents a very different set of assembly challenges compared to the model considered in this work. In addition, some cytokine-receptor binding events induce conformational changes in the receptor that influence receptor dimerization, an effect which is also neglected in our model. As with the

antibody case, most of these structures contain at least one “weak” interaction, but we nonetheless ignore them due to the fact that they do not conform to the assumptions our model. Of the 82 ring structures in the initial dataset, 15 belonged to this class.

3. **Complexes Produced through Proteolysis** This class consists of cases where a subset of the individual chains in the structure are produced when a single chain is cleaved in two. The majority of these cases involve proteases (e.g. trypsin) complexed with protease inhibitors. Proteases such as trypsin are synthesized and fold as a single polypeptide chain (the zymogen). Activation of the zymogen involves a proteolytic cleavage event in which this single chain is cut in two. This produces two polypeptide chains that are intricately folded with one another. In these structures, the interaction between the two chains of the protease is not formed through bimolecular association, but rather through folding as a single chain and subsequent cleavage. This situation is clearly distinct from the assembly dynamics considered in our models, although it is again the case that the protease-inhibitor interactions are considerably weaker than the interaction between the two proteases (a fact which would support our model if the data were included). Of the 82 structures in the initial dataset, 8 belonged to this class.
4. **Miscellaneous** This class consisted of one structure in which binding between two of the subunits was induced by a small molecule (FK506), and one case in which the complex assembles around DNA. Neither case conforms to the assumptions of our model, and so these two structures are also removed from the set.

After curating the 82 structures of heteromeric three-membered rings from 3D Complex, we obtained 29 structures for which the biological system represented by the structure seemed to represent a case similar to that considered in our model. Of these 29 structures, many are enzymes (e.g. glutamine amidotransferase) and many serve regulatory functions (e.g. the complex of the transcription factor NF- κ B with its regulator I κ B α). A list of all structures can be found in a table provided as additional supplementary material.

5.2 Structures for heteromeric four-membered chains

To serve as a contrast to the case of the three-membered rings discussed above, we also considered heteromeric four-membered chains. A heteromeric four-membered chain contains exactly the same number of interactions as a three-membered ring, making it possible to perform a direct comparison between the two types of structures. This comparison is particularly informative due to the fact that optimizing assembly in the case of chains will tend to favor uniformly strong interactions (as demonstrated by the analyses in sections 3.2.3 and 4.3.3).

To assemble this dataset we began with the 104 heteromeric four-membered chains found at the QS-90 level of the 3D Complex hierarchy. We removed structures in the following classes from the data set:

1. **Integral membrane complexes** As with the heteromeric rings discussed above, a number of the chain structures in this data set included one or more proteins that exist as integral membrane proteins in the cell. In general, these structures involved the extracellular

domains of receptors (e.g. the T-cell receptor) complexed with ligands and/or other receptors. As with the three-membered rings, we did not consider the particular evolutionary pressures that might arise in the case of interactions involving integral membrane proteins. Of the 104 chain structures in the initial dataset, 23 belong to this class.

2. **Biological assembly mis-annotated** The 3D Complex database is constructed on the basis of the biological assemblies included in default PDB files [17]. As has been noted elsewhere [18, 19], these complexes are sometimes distinct from the assembly considered biologically relevant in the very manuscript in which the crystal structure at issue is reported. In this case, if the paper reporting a particular structure, or other relevant literature, indicated that the biologically active form of the complex was not a four-membered chain, we removed that structure from our data set. In some cases, we found that authors of the paper reporting the structure cited direct size-exclusion chromatography results indicating that the biologically relevant form of the molecule was a dimer, despite the fact that the “biological assembly” in the PDB was reported as a tetramer. Of the 104 chain structures in the initial dataset, we identified 18 for which there was strong evidence that the biologically relevant form of the complex was not in fact a four-membered chain.
3. **Antibody-antigen complexes** A number of structures in the data set consisted of antibodies binding to various antigens. As with the rings above, we removed these structures from our data set as they represent a case of assembly quite different from that considered here. Of the 104 chain structures in the initial dataset, 8 consisted of antibody-antigen complexes.
4. **Complexes produced through proteolysis** This class consisted of protease and lectin molecules in which at least two of the chains in the final structure are synthesized as a single polypeptide sequence which is later cleaved to give the final, active structure. Since two or more of the chains in these structures do not interact with one another through a bimolecular association event, we did not include these cases in our analysis. Of the 104 chain structures in the initial dataset, 7 involved complexes in which this type of proteolytic cleavage was involved.
5. **DNA-binding complexes** These structures consisted of complexes that assemble around specific DNA sequences. Since assembly on DNA is not considered in our model, we did not include these cases in our analysis. Of the 104 chain structures in the initial dataset, 4 involved complexes assembling on DNA.

After curating the data set into the above classes, we obtained a dataset of 44 heteromeric four-membered chains; as discussed in section 5.3 below, 11 of these structures actually represented four-membered rings upon further analysis. Of the remaining 33 structures, most represent either enzymes or enzyme-inhibitor complexes. A list of these structures is provided as an additional supplementary table.

5.3 Calculating changes in non-polar surface area

Examining affinity distributions in rings and chains using solved protein structures involves estimating the binary binding affinities between components of the structure. Here we focus on measuring the change in Solvent-Accessible Non-Polar Surface Area ($\Delta\text{SASA}_{\text{NP}}$), which has been shown to correlate with binding affinities in some studies [20, 21]. We used the software package POPS [22] to perform this calculation. We proceeded by creating three separate PDB files: file (1) contains only the atoms (ATOM records) that belong to the residues of the first chain (say, chain “A”), file (2) contains only the atoms that belong to the residues of the second chain (“B”), and file (3) contains the atoms from both chains (“A + B”). Next, we calculated the non-polar solvent-accessible surface area for each file separately using POPS. This area is marked as “hydrophobic” in the POPS output. We then calculated $\Delta\text{SASA}_{\text{NP}}$ as the difference between the sum of these areas for each domain separately and the area for the domains combined:

$$\Delta\text{SASA}_{\text{NP}}(\text{A}, \text{B}) = \text{SASA}_{\text{NP}}(\text{A}) + \text{SASA}_{\text{NP}}(\text{B}) - \text{SASA}_{\text{NP}}(\text{A} + \text{B}). \quad (41)$$

In other words, $\Delta\text{SASA}_{\text{NP}}$ is calculated as the SASA_{NP} of file (1) plus the SASA_{NP} of file (2) minus the SASA_{NP} of file (3). We used the definition in equation 41 to calculate $\Delta\text{SASA}_{\text{NP}}$ for every pair of chains in the structure of interest.

For the curated three-membered heteromeric ring structures, all of the cases yielded the expected ring-like topology when subjected to this analysis. The case with the four-membered chains was more complex. Many of the structures in this case actually contained more than 3 interactions (defined as $\Delta\text{SASA}_{\text{NP}}$ greater than some cutoff). Since the only way to include more than 3 unique interactions in a graph of four nodes involves creating a cycle, this analysis indicated that a number of the structures considered as “chains” in 3D Complex actually contained rings. We used the minimum “affinity” observed for rings (129.67 \AA^2) as a cutoff and found that 11 of the 44 curated chains actually exhibited significant ring-like structure, a fact that we confirmed through simple visual analysis of the structure itself in each case. This analysis left 33 “true” four-membered heteromeric chains in our data set.

It is important to note that equation 41 represents a very rough and imperfect measure of affinity [20, 21, 23]. As such, we also explored using total interface area (as opposed to non-polar area), fractional surface area (defined as the area of the interface divided by the entire SASA of the binary complex), and the ΔG values annotated in PISA [24]. We found that these alternative definitions of affinity were strongly correlated with $\Delta\text{SASA}_{\text{NP}}$ ($R^2 > 0.8$ in each case), thus yielding nearly identical results to those discussed for $\Delta\text{SASA}_{\text{NP}}$ below. Of course, all of these definitions are imperfect [23], but in the absence of empirical data regarding affinities in this case, the structural analysis presented here represents the only available test of our predictions.

5.4 Comparing rings and chains

As discussed above, every structure in our ring or chain data sets contains exactly three interactions with $\Delta\text{SASA}_{\text{NP}} > 129.67 \text{ \AA}^2$. For any structure, we can thus define the largest interface (or “Strong” interaction, denoted S), the smallest interface (or “Weak” interaction, denoted W), and the interface in between the two (i.e. the “Medium” interaction, denoted M).

Fig. 4A in the main text represents a summary comparison of the S and W interactions for rings and chains, which we expand upon here.

Figure 24 shows a histogram, the kernel-smoothed density as well as a “rug” plot of the raw data for the S interactions of both rings and chains, while Figure 25 represents the same plot for the W interactions. In the case of S interactions, the mean $\Delta\text{SASA}_{\text{NP}}$ for rings is $\overline{S}_{\text{R}} = 2099 \text{ \AA}^2$, while for chains we have $\overline{S}_{\text{C}} = 1631 \text{ \AA}^2$. To test if this difference in means is significant, we performed a simple random permutation test with 10^5 replicates in the statistical package R [11]. We found that the difference is not significant after correcting for multiple comparisons (the uncorrected $p = 0.0441$). For weak interactions, we have $\overline{W}_{\text{R}} = 531 \text{ \AA}^2$ for rings and $\overline{W}_{\text{C}} = 914 \text{ \AA}^2$ for chains. The difference in means for weak interactions was considerably more significant ($p = 6 \cdot 10^{-5}$). From this analysis we can conclude that the strong interactions in rings are, on average, stronger than those for chains (with the caveat of weak statistical significance), while the weak interactions in rings are considerably weaker than the weak interactions in chains.




Figure 24: Strong interactions in both rings and chains. (A) Here we plot the distribution of “S” affinities for rings. The black boxes represent a histogram of the data, while the red line corresponds to a kernel-smoothed density. The blue lines on the x-axis are a “rug” plot of the data, where each line represents the affinity for a single ring in the data set. The rug plot is provided to give a sense for the data underlying both the histogram and the kernel-smoothed density. (B) A plot as in panel A, but for the S affinities in chains.

It is important to note that the assembly properties of a ring or chain with a particular set of affinities will vary strongly with total monomer concentration (see section 4.1.1). We thus also considered the weak-to-strong interaction ratios (W/S); a plot of the ratio densities for rings and chains can be found in Fig. 4B of the main text. Figure 26 shows this density, as well as histograms and rug plots, for rings and chains separately. Again using a permutation test, we found that the mean ratios for rings, $\overline{(W/S)}_{\text{R}} = 0.309$, is significantly smaller than that for chains $\overline{(W/S)}_{\text{C}} = 0.627$ ($p = 10^{-5}$). Comparison of both the absolute affinity and relative affinity distributions reveals that the weakest interaction in rings is significantly weaker than the weakest interaction in chains, as our assembly models would predict.




Figure 25: Weak interactions in both rings and chains. (A) Here we plot the distribution of “W” affinities for rings. The black boxes represent a histogram of the data, while the red line corresponds to a kernel-smoothed density. The blue lines on the x-axis are a “rug” plot of the data, where each line represents the affinity for a single ring in the data set. The rug plot is provided to give a sense for the data underlying both the histogram and the kernel-smoothed density. (B) A plot as in panel A, but for the W affinities in chains.




Figure 26: The weak-to-strong ratio for both rings and chains. (A) Here we plot the distribution of the weak to strong ratio (W/S) for rings. The black boxes represent a histogram of the data, while the red line corresponds to a kernel-smoothed density. The blue lines on the x-axis are a “rug” plot of the data, where each line represents the affinity for a single ring in the data set. The rug plot is provided to give a sense for the data underlying both the histogram and the kernel-smoothed density. (B) A plot as in panel A, but for the W/S ratios found in chains.

Interestingly, the kernel-smoothed density estimates and the histograms for both rings and chains demonstrate considerable bimodality (see Figure 26). For the rings, we can divide the data into

the first peak (with $(W/S)_R < 0.5$) and the second peak ($(W/S)_R > 0.5$). As can be seen from Figure 26, the majority of rings (24 of 29) belong to the first peak, and taking this peak alone we have $\overline{(W/S)}_R = 0.225$, which is significantly smaller than the average for the entire sample. The 5 points belonging to the second peak have a much higher average, with $\overline{(W/S)}_R = 0.712$. Similarly, we can use a ratio of 0.5 to divide the chains into two peaks, and for 11 structures in the smaller ratio peak we have $\overline{(W/S)}_C = 0.319$ while for the 23 points in the second peak we have $\overline{(W/S)}_C = 0.781$. As Figure 27 demonstrates, normal quantile-quantile plots for the major peaks in both cases (the smaller-ratio peak for the rings and the larger-ratio peak for the chains) reveal that both can be well-approximated by Gaussian distributions (the smaller peaks in both cases contain too few data points to support conclusions regarding normality). Although we do not have enough data to make this point conclusively, we can speculate that both ratio distributions are the result of two combined Gaussian distributions, one with a smaller average ratio, and one with a larger.




Figure 27: Quantile-Quantile plots comparing the ratio distributions of rings and chains with normal distributions. (A) Here we compare the quantiles of the small-ratio W/S distribution for rings (e.g. those rings with $W/S < 0.5$) to the quantiles of the normal distribution. The “Sample Quantiles” on the y-axis are taken from our data on rings, while the “Theoretical Quantiles” for the normal distribution are computed in R. The solid black line represents a linear fit to the Q-Q data. The linear fit in this case is fairly good, with strong statistical significance for both the slope and intercept terms ($p < 2 \cdot 10^{-16}$). Although there are slight systematic deviations from the straight line throughout the range of quantiles, this result indicates that the small-ratio distribution is approximately Gaussian in character. (B) In this case we compare the quantiles of the large-ratio W/S distribution for chains (e.g. those chains with $W/S > 0.5$) to the quantiles of the normal distribution. As in panel A, the “Sample Quantiles” on the y-axis are taken from our data on chains, while the “Theoretical Quantiles” for the normal distribution are computed in R. The solid line is again a linear fit to the Q-Q data; the fit in this case is excellent, with $p < 2 \cdot 10^{-16}$ for both the slope and intercept terms. There is less systematic deviation from the fit in this case, indicating that the large- W/S peak for chains can be fairly well approximated Gaussian distribution.

As can be seen from Fig. 4A in the main text, the smaller-ratio peak for the chains overlaps with that peak for the rings, and vice versa. Using a permutation test, we found that smaller-ratio chains had a significantly higher average than smaller-ratio rings (although the significance is

fairly weak after correcting for multiple comparisons, with an uncorrected $p = 0.0199$), while the averages for the larger-ratio chains and rings were statistically indistinguishable ($p = 0.254$).

We thus find that 24 of the 29 rings considered in our data set conform to our prediction that rings will contain at least one “weak” interaction. The 5 remaining structures do not seem to contain a weak interaction (i.e. they seem to belong to the chain distribution rather than the ring one). These cases may represent situations in which the cell can accommodate sub-optimal assembly efficiency, or cases in which mechanisms not considered in our model (e.g. chaperone-mediated assembly) are involved.

5.5 Gaussian control

Although the above analysis indicates that rings and chains, on average, conform to the expectations of our model, it is unclear to what extent our observations represent evolutionary pressures on affinities. For instance, say we have some positive continuous random variable X that follows an underlying probability density p . Sampling three instances of X from this distribution will naturally result in a “largest” and “smallest” value for the sample, and the ratio of these two numbers will always be less than 1.

We thus conducted a control to test whether we would observe W/S ratio distributions similar to the results for rings and chains when the interactions themselves are sampled from a single underlying distribution. We considered a simple model in which a “structure” is constructed by sampling 3 $\Delta\text{SASA}_{\text{NP}}$ from a Gaussian distribution. We set the mean μ of the Gaussian to be the mean of all the interactions ($A = S \cup M \cup W$) in a given data set; for rings we have $\overline{A_R} = 1255 \text{ \AA}^2$ and for chains we have $\overline{A_C} = 1296 \text{ \AA}^2$. The standard deviation σ of the Gaussian was set to the sample standard deviation in each case: $s(A_R) = 966 \text{ \AA}^2$ and $s(A_C) = 700 \text{ \AA}^2$.

Of course, $\Delta\text{SASA}_{\text{NP}}$ values cannot be negative, and indeed in section 5.3 we defined a cutoff for considering only interactions with $\Delta\text{SASA}_{\text{NP}} > 129.67 \text{ \AA}^2$ as valid. In order to mimic these constraints, we must thus introduce a minimum affinity (129.67 \AA^2) and reject affinities below that cutoff. If we only implement a lower bound, however, the set of sampled interactions from the above procedure will exhibit a mean significantly different from the underlying Gaussian used to construct the data. To prevent this from happening, we implement an upper bound such that the z-score of this upper bound is equal to the absolute value of the z-score for the chosen minimum (that is, $z_{\text{max}} = -z_{\text{min}}$). This allows us to construct a distribution of random affinities from the underlying Gaussian with a minimum possible affinity that is nonetheless symmetric and exhibits the defined average.

Each random structure sampled from the distribution as defined above has a S, M and W interaction, and for each structure we calculate the W/S ratio. A “model” data set is constructed from N such structures, where $N = 29$ for rings and 33 for chains to mimic the distributions we observe in the real data. We constructed 10^4 such data sets (for a total of $2.9 \cdot 10^5$ structures in the case of the rings) and asked what fraction of these random datasets exhibited $\overline{W/S}$ ratios smaller than or equal to that observed for the rings and larger than or equal to that observed for chains.

In every case, we find that this Gaussian control is unlikely to explain the data: $p = 9 \cdot 10^{-4}$ for

rings and $p < 10^{-4}$ for chains.

Fig. 4C in the main text is meant to summarize the results of this control graphically. In that case, we have a single Gaussian distribution with an average taken to be approximately that observed for rings and chains ($\mu = 1250 \text{ \AA}^2$). We vary the standard deviation from $\sigma = 25 \text{ \AA}^2$ to 2500 \AA^2 and maintain a cutoff of 129.67 \AA^2 . In this case, we take $N = 30$ for each data set and we construct 10^4 data sets for each value of σ . In Figure 4C we plot the average W/S and 95% confidence intervals for random data sets constructed this way as a function of σ/μ .

It is important to note that we have controlled here only for one type of underlying distribution; namely a Gaussian with a particular minimum affinity cutoff. Although this control is clearly unlikely to produce the data, one could potentially find some other single underlying distribution of affinities that could. As Fig. 4B in the main text indicates, however, even if this is the case, one could argue that evolution has selected parameters for this underlying distribution (e.g. μ and σ) such that configurations with optimal assembly characteristics are likely.




Figure 28: Relationship between the change in solvent-accessible non-polar surface area and binding free energy. The squares in the plot represent data taken from Table 1 in reference [21]; in this case the ΔG_b values are obtained from experimental measurements, and we determined the $\Delta \text{SASA}_{\text{NP}}$ directly from the corresponding crystal structures using POPS [22] as described in section 5.3 above. The green line is a linear fit to the data, which yields an $R^2 = 0.47$.

5.6 Affinities for the interactions in the crystal structures of rings

One can use available crystal structures of interacting proteins for which affinities are known to investigate the quantitative relationship between $\Delta \text{SASA}_{\text{NP}}$ and ΔG_b^0 . Using a recently-published

data set of 20 such structures [21], we found a roughly linear relationship between the two, with an R^2 of 0.47 (see Figure 28). Although the correlation is imperfect, the linear fit allows us to map $\Delta\text{SASA}_{\text{NP}}$ values into ΔG_b^0 values and thus K_D s for the interactions in our data set. The equation we obtain in this case is:

$$\Delta G_b^0(\text{A}, \text{B}) = -0.015 \cdot \Delta\text{SASA}_{\text{NP}}(\text{A}, \text{B}) - 4.17$$

with $\Delta\text{SASA}_{\text{NP}}(\text{A}, \text{B})$ given in \AA^2 and $\Delta G_b^0(\text{A}, \text{B})$ in kcal mol^{-1} . Assuming approximately room temperatures (i.e. $RT \approx 0.6 \text{ kcal mol}^{-1}$), the average K_D for strong bonds in our rings data set is $8.0 \cdot 10^{-12}$, and the average K_D for weak bonds is $1.8 \cdot 10^{-6}$. Interestingly, these are very close to the values used for Figures 2A and 3B in the main text, as well as the optimum values obtained in our analysis of heteromeric rings (Fig. 3A of the main text). As mentioned above (and as is clear from Figure 28), $\Delta\text{SASA}_{\text{NP}}$ is only a very rough measure of actual binding affinity [20, 21, 23]; these results simply indicate that the K_D values we use for “strong” and “weak” bonds in the text (e.g. 10^{-12} and 10^{-6} M, respectively) are at least broadly consistent with the range of affinities one would expect given the buried surface areas in the crystal structures of homomeric rings.

References

1. Saiz L, Vilar JM (2006) Stochastic dynamics of macromolecular-assembly networks. *Mol Syst Biol* 2:2006 0024.
2. Henriksen NE, Hansen FY (2008) *Theories of Molecular Reaction Dynamics: The Microscopic Foundation of Chemical Kinetics* (Oxford University Press, Oxford).
3. Camacho CJ, Kimura SR, DeLisi C, Vajda S (2000) Kinetics of desolvation-mediated protein-protein binding. *Biophys J* 78:1094–105.
4. Minh DD, et al. (2005) The entropic cost of protein-protein association: a case study on acetylcholinesterase binding to fasciculin-2. *Biophys J* 89:L25–7.
5. Grziwa A, et al. (1994) Dissociation and reconstitution of the thermoplasma proteasome. *Eur J Biochem* 223:1061–7.
6. Nakabayashi J, Sasaki A (2006) A mathematical model for apoptosome assembly: the optimal cytochrome c/apaf-1 ratio. *Journal of Theoretical Biology* 242:280–7.
7. Sharma S, Hoskins JR, Wickner S (2005) Binding and degradation of heterodimeric substrates by clpap and clpxp. *J Biol Chem* 280:5449–55.
8. Moore SD, Baker TA, Sauer RT (2008) Forced extraction of targeted components from complex macromolecular assemblies. *Proc Natl Acad Sci USA* 105:11685–90.
9. The MathWorks, Inc. (2010) *MATLAB Version 7.11* (The MathWorks, Inc., Natick, Massachusetts).
10. Wolfram Research (2008) *Mathematica Version 7* (Wolfram Research, Champaign, Illinois).

11. R Development Core Team (2010) *R: A Language and Environment for Statistical Computing* (The R Foundation for Statistical Computing, Vienna, Austria).
12. Kress W, Mutschler H, Weber-Ban E (2007) Assembly pathway of an aaa+ protein: Tracking clpa and clpap complex formation in real time. *Biochemistry* 46:6183–6193.
13. Ghaemmaghami S, et al. (2003) Global analysis of protein expression in yeast. *Nature* 425:737–741.
14. Belle A, Tanay A, Bitincka L, Shamir R, O’Shea EK (2006) Quantification of protein half-lives in the budding yeast proteome. *Proc Natl Acad Sci USA* 103:13004–9.
15. Yin P, Choi HMT, Calvert CR, Pierce NA (2008) Programming biomolecular self-assembly pathways. *Nature* 451:318–22.
16. Elowitz MB, Levine AJ, Siggia ED, Swain PS (2002) Stochastic gene expression in a single cell. *Science* 297:1183–6.
17. Levy ED, Pereira-Leal JB, Chothia C, Teichmann SA (2006) 3d complex: a structural classification of protein complexes. *PLoS Comput Biol* 2:e155.
18. Xu Q, et al. (2008) Statistical analysis of interface similarity in crystals of homologous proteins. *Journal of molecular biology* 381:487–507.
19. Xu Q, Dunbrack RL (2011) The protein common interface database (procid)—a comprehensive database of interactions of homologous proteins in multiple crystal forms. *Nucleic Acids Research* 39:D761–70.
20. Horton N, Lewis M (1992) Calculation of the free energy of association for protein complexes. *Protein Sci* 1:169–81.
21. Bougouffa S, Warwicker J (2008) Volume-based solvation models out-perform area-based models in combined studies of wild-type and mutated protein-protein interfaces. *BMC bioinformatics* 9:448.
22. Cavallo L, Kleinjung J, Fraternali F (2003) Pops: A fast algorithm for solvent accessible surface areas at atomic and residue level. *Nucleic Acids Research* 31:3364–6.
23. Kastiris PL, et al. (2011) A structure-based benchmark for protein-protein binding affinity. *Protein Sci* 20:482–491.
24. Krissinel E, Henrick K (2007) Inference of macromolecular assemblies from crystalline state. *Journal of molecular biology* 372:774–97.

B

The noise is the signal: information flow in single cells and cellular populations

AUTHORSHIP CONTRIBUTION

E. Deeds and R. Suderman proposed the relationship between single-cell noise and population-level response; I designed and performed the experiments aimed at exploring this relationship in the context of apoptosis. In addition I reviewed, tested, and documented the code for calculating channel capacity (contributions on GitHub at <https://github.com/ryants/EstCC/graphs/contributors>). I prepared several figures related to the experimental data and contributed to the writing of the manuscript.

ABSTRACT

Signal transduction networks allow cells to make decisions based on their intracellular state and information in the environment. Cheong *et al.* recently demonstrated that noise significantly diminishes the

fidelity of signaling in eukaryotic cells: the majority of the networks they examined transmitted less than 1 bit of information. It is unclear how networks that control critical cell fate decisions (e.g. cell division and apoptosis) can function effectively with such low levels of information transfer. For apoptosis in human cells, we demonstrate the existence of an inherent trade-off between the amount of information transferred in individual cells and the information available to control population-level responses; noise that reduces information transfer to less than one bit at the single-cell level allows 3–4 bits of information to be transmitted at the population level. For processes such as eukaryotic chemotaxis, in which single cells are the key biological unit, we find high levels of information transmission at a single-cell level. Thus, the low levels of information transfer observed previously are unlikely to represent an inherent physical limit. Instead, we propose that signaling networks can exploit noise to maximize population-level information transfer. This is particularly critical for discrete changes in cell fate (e.g. apoptosis) in which the key variable is the fraction of cells engaged. Our findings provide a conceptual framework for rationalizing the high levels of noise in metazoan signaling networks and have important implications for the development of drugs that target these networks in the treatment of cancer and other diseases.

SIGNIFICANCE STATEMENT

Cells rely on signaling networks to detect changes in their environment and respond appropriately. Recent evidence suggests that there are high levels of *noise* in eukaryotic signaling, and it is currently unclear how cells can make critical cell fate decisions (e.g. programmed cell death) on the basis of low levels of information about their environment. Here, we show that high noise levels are actually critical when a system needs to control the behavior of *populations* of cells. In contrast, when the key biological unit is a single cell, we show that the impact of noise on signaling is much less pronounced. Understanding the regulation of noise will be key to developing a complete picture of information processing in cells.

INTRODUCTION

Signaling networks allow cells to sense intra- and extra-cellular concentrations of cytokines, nutrients, ions, etc., and execute both discrete and continuous changes in cell state in response to those signals

(Spencer et al., 2009; Chen et al., 2012; Balázsi et al., 2011; Suderman and Deeds, 2013; Rowland et al., 2012). Apoptosis and commitment to cell division are typical of binary responses, whereas directed cell movement and induced gene expression are typical of continuously variable responses (Chen et al., 2012; Suderman and Deeds, 2013; Bashor et al., 2008a; Bardwell, 2004). Dysregulation of intracellular signaling has been implicated in a wide range of diseases including cancer, chronic inflammation, neurodegeneration, etc. (Sebolt-Leopold and Herrera, 2004), and developing a fundamental understanding of cellular information processing is instrumental in developing rational strategies aimed at treating those diseases (Rowland et al., 2012; Fallahi-Sichani et al., 2013). While signaling networks have been the subject of intense experimental and theoretical study for decades, it has only recently become possible to measure the response to signals at the level of individual cells (Spencer et al., 2009; Cheong et al., 2011; Selimkhanov et al., 2014). These studies have revealed that signaling networks are subject to significant *noise*, which manifests itself within single cells as stochastic fluctuations in the activities of signaling proteins and as cell-to-cell variability within genetically identical cell populations (Spencer et al., 2009; Balázsi et al., 2011; Fallahi-Sichani et al., 2013; Flusberg et al., 2013; Albeck et al., 2008; Feinerman et al., 2008; Kaern et al., 2005; Ahrends et al., 2014; Shalek et al., 2014; Mehta et al., 2009). While the presence of high levels of noise in human signaling networks seems ubiquitous, the ultimate physiological role of this heterogeneity remains unclear.

Information theory (Shannon, 1948) provides a powerful analytical framework for quantifying the impact of noise on the ability of a system to transmit information. Levchenko and co-workers pioneered the application of information theory to signaling in mammalian cells (Cheong et al., 2011) with the concentration of an extracellular ligand (e.g., the inflammatory cytokine TNF- α) serving as the input to a (potentially noisy) intracellular signaling network (or channel), ultimately leading to a downstream response that can be experimentally measured (e.g. the nuclear translocation of NF- κ B). The information carried by the channel is quantified by the mutual information, I :

$$I(X; Y) = \sum_x \sum_y p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \quad (\text{B.1})$$

where X is the probability distribution of the signal and Y is the distribution of the response (Cheong

et al., 2011; Shannon, 1948). The base of the logarithm determines the units of the mutual information: the conventional base 2 quantifies information in “bits.” Since the value of I depends on the input distribution, the mutual information of a signaling channel represents a combination of the properties of the signal and the intrinsic limits of the channel itself. It is thus problematic to use I as a measure for evaluating or comparing information flow in various signaling networks, since the *in vivo* distribution of signal values is rarely known. As a result, it is often more informative to focus on the maximum possible information that a channel can carry, or the channel capacity, C :

$$C = \sup_{p_X(x)} I(X; Y) \quad (\text{B.2})$$

where the supremum (the least upper bound) is evaluated over all possible choices of the probability distribution of the input. The channel capacity is an inherent feature of the channel: the larger the value, the more information that a channel can theoretically transmit (Cheong et al., 2011; Shannon, 1948). In the case of cellular responses to TNF and other cytokines, Cheong et al. found that the channel capacity is generally less than 1 bit for molecular responses at the single-cell level (these values are summarized in Table B.1, entries 1-4, 9) (Cheong et al., 2011; Bao et al., 2010; Cohen-Saidon et al., 2009; Coppey et al., 2008). Since the number of distinct signal values that can be resolved is 2^C , the implication is that many intracellular signaling networks cannot reliably distinguish between the presence or absence of TNF, EGF, and other signaling molecules ($C < 1$ bit, Table B.1) (Cheong et al., 2011). More recent work has focused on characterizing various strategies that cells might employ to achieve higher levels of information transfer. For instance, Lee *et al.* demonstrated that mechanisms such as fold-change detection (in which cells are sensitive to the ratio between a steady-state and induced signal) decrease the impact of noise on the propagation of TNF-induced signals (Lee et al., 2014). As described below, however, we found that the channel capacity between TNF concentration and the downstream transcriptional response remains below 1 bit despite the use of fold-change detection in this system (entry 5, Table B.1). Wollman and co-workers recently demonstrated that using multiple time points from the trajectory of a molecular response (e.g., Erk activation over time) can significantly increase channel capacities. While this dynamic approach to information flow clearly can increase C , it is currently unclear how cells might actually implement this

kind of mechanism at the molecular level [Selimkhanov et al., 2014](#).

Signal (molecular)	Response (molecular)	C (bits)	Data	Calculation
1. TNF	NF- κ B	0.92 ± 0.01	Cheong <i>et al.</i>	Cheong <i>et al.</i>
2. PDGF	NF- κ B	0.67 ± 0.01	Cheong <i>et al.</i>	Cheong <i>et al.</i>
3. EGF	Erk (fold-change)	0.60 ± 0.03	Cohen-Saidon <i>et al.</i>	Cheong <i>et al.</i>
4. UDP	Peak Ca^{2+}	1.22 ± 0.03	Bao <i>et al.</i>	Cheong <i>et al.</i>
5. TNF	A20 transcripts	0.62 ± 0.08	Lee <i>et al.</i>	this work
6. TRAIL	Casp-8 activity	0.66 ± 0.02	this work	this work
7. TRAIL	Casp-3 activity	0.33 ± 0.01	this work	this work
8. α -factor	<i>pFUS1</i> -GFP	2.26 ± 0.05	Bashor <i>et al.</i>	this work
Signal (position)	Response (molecular)	C (bits)	Data	Calculation
9. Embryo perimeter	Phosphorylated Erk	1.61 ± 0.05	Copey <i>et al.</i>	Cheong <i>et al.</i>
Signal (position)	Response (motion)	C (bits)	Data	Calculation
10. Bacterium	neutrophil motion	2.14 ± 0.22		this work
11. cAMP	<i>Dictyostelium</i> motion	2.24 ± 0.07	Firtel Lab	this work
Signal (molecular)	Response (population)	C (bits)	Data	Calculation
12. TRAIL	% dead (HeLa cells)	3.46 ± 0.03	this work	this work
13. TRAIL	% dead (MCF10A cells)	3.36 ± 0.03	this work	this work

Table B.1: Calculated channel capacities for experimental data. The channel capacity for population-level response in HeLa cells was calculated using 1000 cells per TRAIL concentration and all population-level channel capacities were calculated using 100 independent populations.

One limitation of previous work is the focus on measuring the activity of signaling intermediates, such as nuclear localization of the NF- κ B transcription factor, rather than a cellular phenotype as an output ([Cheong et al., 2011](#); [Lee et al., 2014](#)). This makes it difficult to interpret the functional significance of low channel capacities. We therefore focused our analysis on an unambiguous terminal phenotype: life or death as regulated by TNF-Related Apoptosis-Inducing Ligand (TRAIL). TRAIL induces apoptosis by binding to cell surface receptors, initiating formation of *death-inducing signaling complexes* or DISCs. These complexes then activate initiator caspases (ICs), starting a sequence of biochemical events resulting in *mitochondrial outer membrane permeabilization* (MOMP). The release of numerous mitochondrial proteins into the cytosol and subsequent formation of the apoptosome then promotes activation of the effector caspases (ECs), ultimately leading to cell death (Figure [B.1A](#)). Dramatic cell-to-cell variability has

been observed in the responses of clonal cell cultures to TRAIL (and other death ligands): whereas a subset of cells dies within 2-8 hr of ligand exposure, others survive indefinitely. When these survivors are re-assayed for TRAIL sensitivity following outgrowth, the same fractional killing is observed, showing that variability is a stable property of the cell population. Molecular studies have shown that this variability arises from extrinsic noise in receptor-to-caspase signaling networks (Spencer et al., 2009; Flusberg et al., 2013; Albeck et al., 2008).

We found that the channel capacity between TRAIL dose and IC or EC activity (measured at the single-cell level) is significantly less than 1 bit, similar to other molecular responses to cytokine signals (Table B.1, entries 6 and 7). Interestingly, however, we found that the channel capacity between TRAIL dose and the *fraction of cells that die* at that dose is much higher, around 3-4 bits. Based on these findings, we developed a simple mathematical model that allowed us to characterize the fundamental trade-off between the amount of information individual cells can have about their environment and the amount of information that can be used to control decision-making at the population level. We also demonstrated that the low channel capacities generally observed for single-cell responses are not a result of inherent biophysical limitations: by analyzing data on eukaryotic chemotaxis and mating for yeast cells, we found that some signaling networks are capable of transmitting well over 2 bits of information to individual cells. Ultimately, our work suggests that noise in cell signaling is likely highly regulated. When the key physiological output is the behavior of a single cell (as in chemotaxis or mating), noise is likely suppressed (Selimkhanov et al., 2014; Lee et al., 2014) to enable high levels of information transfer to those cells. When the key physiological output is the fraction of cells in a tissue or population that undertake a certain decision (e.g. commitment to apoptosis or cell division), noise is likely exploited so that information can be transferred at the population level.

RESULTS

INDIVIDUAL CELLS RESPONDING TO TRAIL EXHIBIT A LOW CHANNEL CAPACITY

To measure the channel capacity of the extrinsic apoptosis signaling cascade (Figure B.1A), HeLa cells were treated with TRAIL for 11 hr over a range of ligand concentrations from sub- to super-physiological, and




Figure B.1: Cell-to-cell variability in response to a range of TRAIL doses. **(A)** TRAIL activates the extrinsic apoptosis signaling pathway through activation of the initiator caspase (IC) Casp-8 via death-inducing signaling complexes (DISCs). Active Casp-8 then activates the effector caspase (EC) Casp-3, via two mechanisms: direct cleavage and mitochondrial outer membrane permeabilization or MOMP, which induces formation of the apoptosome, another activator of Casp-3 (13). **(B)** Measurement of cleaved EC and IC substrates by flow cytometry (13) show that HeLa cells have a highly variable response to TRAIL across a wide range of doses ($n = 60,000$ cells per TRAIL dose). The solid line is the minimum density in the bimodal EC response (~ 2.8 in \log_{10} units) and acts as a threshold for apoptosis, whereas the dashed line marks the average IC response for non-apoptotic cells. **(C & D)** We used kernel density estimators in the R statistical software package (*R: A Language and Environment for Statistical Computing 2010*) to estimate TRAIL-dependent response distributions for IC (C) and EC (D) activity. The fraction of EC activity above the threshold is proportional to the number of apoptotic cells (13) indicating that approximately 50% of cells survive the maximum TRAIL dose.

molecular responses in single cells were measured by flow cytometry [Albeck et al., 2008](#). The level of cleaved caspase-3 (cC3) served as a measure of the time-integrated activity of receptor-proximal ICs and cleaved PARP (cPARP) served as a measure of downstream EC activity (Figure [B.1A](#)). Previous studies have shown that TRAIL exposure results in a dose-dependent increase in IC activity that varies significantly from cell-to-cell; in any single cell, when IC activity exceeds a threshold set by anti-apoptotic Bcl-2 proteins, ECs are activated and the cell proceeds inexorably to death (Figures [B.1B-D](#)) ([Spencer et al., 2009](#); [Albeck et al., 2008](#)).

While Eqs. [B.1](#) & [B.2](#) seem concise at first glance, estimation of the mutual information and channel capacity is a nontrivial challenge, and numerous approaches have been proposed and implemented ([Steuer et al., 2002](#); [Paninski, 2006](#); [Kraskov et al., 2004](#)). In order to facilitate comparison between our calculations and those performed by Cheong *et al.*, we designed a software package to estimate mutual information based on the binning procedure they applied in their work (see Materials and Methods, and the Supporting Information for further details) ([Cheong et al., 2011](#); [Nemenman et al., 2004](#)). This software is freely available as an open-source project (<https://github.com/ryants/EstCC>). Using this software and the distributions of IC and EC activity in single cells, we calculated a channel capacity between TRAIL dose and IC activity of $C \sim 0.66$ bits and between TRAIL and EC activity of $C \sim 0.33$ bits (entries 6 and 7, Table [B.1](#)). We observed similar values for TRAIL to IC channel capacity in surviving cells (i.e. those with low cPARP levels, see Supporting Information). As an additional control, we calculated the channel capacity between IC activity and both EC activity and cell fate obtaining a C of 1.23 and 0.85 bits, respectively (see Supporting Information). This confirms that measured IC activity is a relevant intermediate signal for extrinsic apoptosis (since it contains almost all the information necessary to specify the binary cell-fate decision), and that noise accumulates progressively through the apoptosis signaling pathway ([Spencer et al., 2009](#)). High levels of noise and low levels of information transfer are thus a feature of the TRAIL network across multiple biologically relevant measures.

POPULATIONS OF CELLS RESPONDING TO TRAIL EXHIBIT HIGH CHANNEL CAPACITY

However, when we examined the channel capacity between TRAIL dose and phenotypic response at the population level we obtained a very different result. The combination of noise and a threshold can allow a fraction of cells in a population to make a discrete decision in response to a signal (Spencer et al., 2009; Feinerman et al., 2008; Kaern et al., 2005; Ahrends et al., 2014; Shalek et al., 2014; Miller et al., 2012). For either of two cell types (transformed HeLa and non-transformed MCF10a cells), we found that the fraction of cells surviving exposure to TRAIL gradually decreased as the concentration of ligand increased over a 10^3 -fold range (Figure B.2). The fraction of cells dying at any given TRAIL dose in both experiments showed comparatively little variance between replicate experiments (Figure B.2). As a result of this relatively low variability, the channel capacity between TRAIL dose and the fraction of cells undergoing apoptosis was much higher than what we observed for the molecular response in single cells, between 3.4 and 4 bits depending on the population size (entries 12 and 13 in Table B.1, and Figure B.3 below). Since the variability between dose-response replicates is at least partly technical in nature (e.g., due to handling of the cells during the experiment), these values represent lower bound estimates of the true biological channel capacity.

UNDERSTANDING THE TRADE OFF BETWEEN SINGLE-CELL AND POPULATION-LEVEL INFORMATION TRANSFER

To better understand how single-cell noise contributes to high channel capacity at a population level (Balázsi et al., 2011; Feinerman et al., 2008; Kaern et al., 2005; Ahrends et al., 2014; Shalek et al., 2014), we created a simplified model for TRAIL signaling in which signal (S) and response (R) are related by a Hill function modified to account for noise:

$$R = (R_{\max} - R_{\min}) \cdot \frac{S^n}{S^n + K^n} + R_{\min} + \varepsilon \quad (\text{B.3})$$

where K is the concentration of an input ligand that results in a half-maximal response, n is the Hill exponent (a measure of dose-response ultrasensitivity), R_{\min} and R_{\max} represent the range of average re-




Figure B.2: Population-level dose-response relationship for TRAIL-mediated apoptosis **(A)** We used the threshold described in Fig. 1B to map data from HeLa cells to fractional survival at varying TRAIL doses. We recorded a maximal effect at [TRAIL] = 1000 ng/mL (indicated by the dashed line); higher doses of TRAIL lead to less fractional killing in a “ligand squelching” effect that we have consistently observed for this system. Since the channel capacity represents a supremum over all possible probability distributions of input signals, we removed the final point ([TRAIL] = 2000 ng/mL) from our analysis without loss of generality. Error bars indicate sample standard deviation across 3 replicates of 20,000 cells each. **(B)** Fraction of MCF10A cells surviving TRAIL treatment as assayed by methylene blue staining (see supplementary online text) (12) show a graded response similar to that of the HeLa population in (A). Calculation of the channel capacity based on this data yielded $C \sim 3.4$ bits, similar to the value for the HeLa cells (Table B.1).

sponses, and ε is a noise term sampled from a Gaussian distribution with mean $\mu = 0$ and variable standard deviation, σ (Albeck et al., 2008). Estimation of the single-cell channel capacity in this model involves generation of a dose-response data set with N independent cells per M distinct concentrations of input signal, resulting in $N \times M$ ordered (S, R) pairs. By varying σ , we created dose-response data for specified levels of noise (Figure B.3A) (Albeck et al., 2008). To simulate responses at the population level (e.g. the fraction dead), one must map individual cell responses to a discrete phenotype. Our model therefore assumes that individual cells exhibit a phenotypic response (e.g. cell death) when R exceeds a threshold value, (analogous to the threshold set by anti-apoptotic Bcl-2 proteins in apoptosis) (Albeck et al., 2008). At any given signal value we thus have the distribution of R values in a simulated population, and also the fraction of that simulated population that exhibit the phenotypic response.

For simulated populations of $N = 10^2$ to 10^4 cells, the model revealed a striking trade-off between the channel capacity for single cells and for cell populations. When noise is low, the response of individual cells is essentially deterministic, corresponding to a step-like change in the fraction of cells that die as S increases (Figures B.3A and B, blue). At higher levels of noise, the response of individual cells obviously becomes more variable, but this corresponds to a gradual decrease in the fraction of “surviving” cells (Figures B.3A and B, black and red) (Spencer et al., 2009; Feinerman et al., 2008; Kaern et al., 2005; Ahrends et al., 2014; Shalek et al., 2014; Miller et al., 2012). As observed in the experimental data (Figure B.2), the fraction of cells that respond at any given value of S displays relatively low variance between populations of the same size (corresponding to relatively small error bars on the red and black curves in Figure B.3B).

Low noise thus leads to high channel capacities between S and R measured at the single-cell level (over 5 bits), but very low channel capacities between S and the fraction of cells that die (~ 1 bit, Figure B.3C). As the level of noise increases, channel capacity at the single-cell level drops rapidly but at the population level it rises significantly, before falling again (Figure B.3C). The level of noise that optimizes population-level channel capacity varies with the number of cells. For example, with $N = 10^2$ a maximum of $C \sim 2.75$ bits is achieved with $\sigma \sim 1$; with 10^4 cells $C \sim 4$ with $\sigma \sim 2.5$ (Figure B.3D). It should be noted that the precise position of this maximum depends on the spacing of the signal values S : since lower values of sigma correspond to a narrower population-level dose response (Figure B.3B), re-sampling S values more finely




Figure B.3: Relationship between single cell and population-level channel capacity. **(A)** Single-cell dose-response behavior in the initial model described in Eq. B.3. The mean response and sample standard deviation of 1000 independent simulated "cells" is shown for various noise values, relative to a cell death threshold (dashed line). **(B)** Population dose-response behavior from $P = 100$ independent populations with $N = 1000$ cells per signal each. Individual cells' response map to either death or survival according to the threshold in (A); points correspond to the mean and sample standard deviation of the fraction of surviving cells. **(C)** A trade-off exists between single-cell and population-level channel capacity. Increasing noise decreases information transmission in single cells and simultaneously increases the population-level channel capacity up to an optimal noise value. **(D)** Channel capacity in our model correlates with population size in the presence of noise (green, black, red, blue). Additionally, the level of noise needed to maximize channel capacity changes as the population size grows, from low values ($\sigma \sim 1$ for $N = 100$) to higher values ($\sigma \sim 2.5-5$ for $N = 10,000$). Experimental population-level channel capacities (orange) were calculated by taking 100 random subsamples from the set of 60,000 HeLa cells for a range of population sizes.

within the transition region tends to decrease the value of noise that maximizes C (see the Supporting Information). However, so long as there is some minimum spacing between the discrete S values to which a population can be exposed, or any error in generating precise values of S (e.g., experimental error in preparing a precise TRAIL concentration), the maximum observed in Figure B.3D occurs at standard deviations significantly larger than 0.

To examine the effect of population size in our experimental data, we randomly sampled subpopulations of HeLa cells from the total of 60,000 per TRAIL dose that we measured. This revealed a similar dependence of population-level information transfer on population size in our experimental data (Figure B.3D, orange). Taken together, our work demonstrates that the combination of a noisy signaling network with concomitantly low information transfer (Figure B.1 and Table B.1) and a threshold in initiator caspase activity (Albeck et al., 2008) leads to robust information transfer at the level of cell populations (Figures B.2 and B.3).

LOW CHANNEL CAPACITIES OBSERVED PREVIOUSLY LIKELY DO NOT REPRESENT INTRINSIC BIO-PHYSICAL LIMITS

Although noise may ultimately support information transfer to cell populations, it is unclear if the phenomenon discussed above represents cells simply taking advantage of the *inherent noise* in signaling systems, or if noise can be tuned up and down to favor fidelity in either single cells or population-level decisions. To explore this latter possibility, we considered two cases in which individual cells (rather than populations) must respond accurately to environmental stimuli. During *S. cerevisiae* mating, haploid a and α cells must determine if a suitable mating partner is sufficiently close for conjugation to be successful. Cells sense the local concentration of the mating pheromone α factor via a G protein-coupled cell surface receptor and a downstream MAP kinase signaling cascade; when a suitable partner is available for conjugation they reorient their cytoskeletons and initiate a complex transcriptional program (Suderman and Deeds, 2013; Bashor et al., 2008b; Bardwell, 2004). Since the decision to mate results in cell-cycle arrest, we would expect there would be an evolutionary pressure for individual yeast cells to have a relatively high level of information about the availability of mating partners in their environment. Using

the single-cell data from Bashor *et al.* (6), we calculated $C \sim 2.26$ bits between α -factor dose and the transcriptional output of the signaling network, measured by a fluorescent reporter (entry 8, Table B.1). This particular network thus demonstrates a much higher level of fidelity than has been observed for molecular responses to cytokines in metazoan systems (Table B.1).

Another example of a situation in which individual cells are the key biological actors is eukaryotic chemotaxis. We therefore analyzed a classic movie of a human neutrophil “hunting” a bacterial cell, and a movie of a single *Dictyostelium* cell responding to cAMP emanating from a micropipette (both movies are available as supplemental files) (Janetopoulos and Firtel, 2008). Because migrating cells are polar, it is possible to define a cell-based coordinate system using standard tracking software (*CellTrack*) (Sacan *et al.*, 2008). Like others working on distributions of directional movement (Burov *et al.*, 2013), we defined the input as the angle between the chemoattractant (bacterium or micropipette) and the cell axis and the output as the angle of the cell’s subsequent motion (Figure B.4A, Materials and Methods). For both the neutrophil (a representative trajectory is shown in Figure B.4B) and *Dictyostelium* we computed $C > 2$ bits (entries 10 and 11, Table B.1), which is almost certainly a lower bound given that we are simplifying a 3D problem as a 2D search (see Discussion). From these data we conclude that signaling networks in single cells can encode more than 2 bits of information (possibly much more) demonstrating that previous observations of $C \sim 1$ are likely not due to the fact that the inherent noise in biochemical reaction networks limits channel capacities to below 1 bit.

DISCUSSION

Our findings touch on two distinct and complementary aspects of information transfer in signal transduction: single-cell and population-level information processing. In the case of regulatory networks that control apoptosis, the key physiological variable is the fraction of cells responding at a given dose (Ahrends *et al.*, 2014; Miller *et al.*, 2012). In this case, low channel capacity at a single-cell level ($C < 1$) is a corollary of high capacity at a population level ($C \sim 3$ to 4). Said another way, achieving effective control over fractional responses requires a significant heterogeneity at the single-cell level (Spencer *et al.*, 2009; Chen *et al.*, 2012; Feinerman *et al.*, 2008; Kaern *et al.*, 2005; Ahrends *et al.*, 2014; Shalek *et al.*, 2014; Miller *et al.*,




Figure B.4: Schematic for spatial channel capacity calculation. **(A)** A representative trajectory from the neutrophil movie. Points are centers of mass for the bacterium (signal source, red) and neutrophil (motile cell, black). The bold, outlined areas show the cells' perimeters in the trajectory's first frame. Clearly seen here is the bacterium's stochastic random walk-like motion and the neutrophil's smoother tracking of the resulting gradient. **(B)** Black circles represent the motile cell and red circles represent the signal source (either a micropipette or bacterium). Since cells do not instantaneously detect or respond to extracellular stimuli, filled, solid, and dashed circles represent the motile cell and signal source at initial (t), signal-delayed ($t + \Delta t_1$), and response-delayed ($t + \Delta t_1 + \Delta t_2$) time, respectively. We can then calculate the mutual information between the signal (\mathcal{I}_1) and response (\mathcal{I}_2) angles.

2012), and thus low channel capacities when the responses of those cells are assayed at a single-cell level. Since many cytokines regulate population-level behaviors (e.g. control over neural progenitor cell proliferation and differentiation by EGF/NGF (Chen et al., 2012)), it is perhaps not surprising that channel capacities less than one bit have been observed in those cases (Table B.1).

In contrasting cases where individual cells must precisely resolve signals to make decisions in a continuous response space (e.g. finding a mating partner, following gradients or hunting pathogens) we find that the single-cell channel capacity is generally significantly higher than has hitherto been observed ($C \sim 2$, Table B.1) (Figure B.4). As mentioned above, we expect that this value is likely a lower bound, since our chemotaxis data is essentially a 2-dimensional projection of a process that often occurs in 3-dimensional space. Additionally, the data for the Dictyostelium calculation ($C \sim 2.24$, Table B.1, entry 11) exhibits tight distributions around approximately 6 input angles, producing a maximum input entropy (and thus an upper limit on channel capacity) of slightly less than $\log_2(6) = 2.6$ bits. Since the estimated channel capacity is so close to this limit (>85%), we suspect that the calculated spatial channel capacity in this case would increase upon further sampling of signaling space (i.e. the relative angle between the source of the signal and the cell).

The clear conclusion from these findings is that low channel capacity at a single cell level ($C < 1$) does not reflect an inherent limit in the biochemistry of signal transduction, but rather a natural trade-off between the knowledge that individual cells have about their environment and the ability of multicellular organisms to control responses reliably at the population level. With respect to noise levels in these systems, two nonexclusive possibilities exist. The first is that networks that control cellular populations simply exploit noise that arises from stochastic fluctuations in transcription, protein synthesis and related processes whereas chemotactic networks have evolved to suppress it. The second is that some signaling networks have actually evolved higher levels of noise than the underlying biophysics dictates (Shahrezaei and Swain, 2008; Friedman et al., 2006; Cai et al., 2006; Kepler and Elston, 2001). In either case, the physiological importance of noise may explain why drugs that target cellular decision networks have difficulty eliciting complete population-level responses (Fallahi-Sichani et al., 2013). Understanding and ultimately exploiting biological noise is likely to be as important for therapy as it is for metazoan signaling.

MATERIALS AND METHODS

EXPERIMENTAL METHODS

HeLa cells were maintained in DMEM medium (Corning 10-013-CV) with 10% fetal bovine serum and 1% penicillin/streptomycin solution (Life Technologies 15140-122). For TRAIL dose-response assays, HeLa cells were plated at a density of 250k cells/well in 12-well plates (Sigma SIAL0513), allowed to adhere overnight, and treated with varying doses of SuperKiller TRAIL (Axxora ALX-201-115) for 11 hours. Three replicate wells were used for each TRAIL dose to establish the technical variability of the assay. After treatment, medium containing dead cells was transferred to flow cytometry tubes (BD Falcon 352235) containing 2 ml FACS buffer (PBS + 10% fetal bovine serum); cells remaining in the wells were removed by trypsinization, added to the corresponding tubes, pelleted by centrifugation and fixed in 4% paraformaldehyde for 30 minutes. After fixation cells were washed twice in PBS and permeabilized in 100% methanol overnight at -20°C. Cells were stained with primary antibodies to cleaved caspase 3 (rabbit anti-cleaved caspase 3, BD 559565) and cleaved PARP (mouse anti-cleaved PARP, BD 552596) 1:250 in FACS buffer (PBS + 0.1% Tween-20) for 1 hour at 25°C. Cells were washed twice in PBS-T, then treated with secondary antibodies: Alexa-488 donkey anti-rabbit IgG (Life Technologies A-21206) and Alexa-594 donkey anti-mouse IgG (Life Technologies A21203), 1:500 in FACS buffer for 1 hr at 25°C. Cells were washed in PBS-T, resuspended in PBS, and counted on a flow cytometer (BD LSRII), with 20,000 cells analyzed per experimental replicate.

MCF10A cells were obtained from J. Brugge (Harvard Medical School, Boston, MA) and cultured as described (Debnath et al., 2003). For TRAIL dose response assays, MCF10A cells were plated in 96-well plates (Corning 353072) and treated with varying doses of SuperKiller TRAIL for 11 hours. After treatment the cells were washed with PBS and the density of viable cells was assayed by methylene blue staining as described previously (Flusberg et al., 2013)

CALCULATING MUTUAL INFORMATION

The code used to calculate the channel capacity was based primarily on the description of mutual information estimation in Cheong *et al.*'s supplementary texts (Cheong *et al.*, 2011) but was modified in a few ways. Instead of calculating the average mutual information of the “plateau” region of bins, we take the maximum mutual information such that at least one mutual information estimate from 10 randomized data sets includes the value 0.0 in its 95% confidence interval about the mean. The parameters used in our jackknife sampling procedure were also slightly different for our linear function estimation as we implemented a minimum subsampling of 80% of the data instead of the 60% used by Cheong *et al.* (Cheong *et al.*, 2011). We used this method for all mutual information calculations performed in this work in order to ensure accurate comparison between values. The source code can be found at <http://github.com/ryants/EstCC> and a complete description of the estimation procedure can be found in the Supporting Information.

MODEL CONSTRUCTION

All calculations involving the model seen in Eq. B.3 were performed with the following (arbitrarily chosen) parameters: $K = 10$, $n = 6$, $R_{\max} = 30$, and $R_{\min} = 20$. Our range of 20 signal values was chosen such that the minimum and maximum response values in our data set were 10% above and 10% below R_{\min} and R_{\max} , respectively, and the remaining 18 values were evenly distributed in between. In this way, the Hill coefficient governing the slope (or ultrasensitivity) of the response, and the sampled signal space, minimally impacts the channel capacity calculation (see Supporting Information). The threshold value was chosen so that half of the signal values produce an average response below the threshold and half produce an average response above the threshold. In the absence of noise, this selection would result in a channel capacity of 1 bit.

Spatial channel capacity calculation: From the 2D data provided by the CellTrack program, we calculated the mutual information between the initial angle created from the motile cell (ϑ_1) and the signal source and the resulting angle of motion of the motile cell (ϑ_2) as mentioned in the main text (Figure B.4). Since information transmission does not occur instantaneously, we introduced two time-delay fac-

tors: Δt_1 , which is the time necessary for the motile cell to detect the signal, and Δt_2 , which is the time required for the neutrophil to respond to extracellular information. We calculated $C(\vartheta_1, \vartheta_2)$ for a range of $(\Delta t_1, \Delta t_2)$ pairs and reported the maximal value in Table B.1 (values for other sets of time parameters and other calculation methods are recorded in the Supporting Information).

References

- Ahrends, R., A. Ota, K. M. Kovary, T. Kudo, B. O. Park, and M. N. Teruel (2014). “Controlling low rates of cell differentiation through noise and ultrahigh feedback.” *Science (New York, NY)* 344.6190 (June 2014), pp. 1384–1389.
- Albeck, J. G., J. M. Burke, B. B. Aldridge, M. Zhang, D. A. Lauffenburger, and P. K. Sorger (2008). “Quantitative analysis of pathways controlling extrinsic apoptosis in single cells”. *Molecular Cell* 30.1 (Apr. 2008), pp. 11–25.
- Balázsi, G., A. van Oudenaarden, and J. J. Collins (2011). “Cellular decision making and biological noise: from microbes to mammals.” *Cell* 144.6 (Mar. 2011), pp. 910–925.
- Bao, X. R., I. D. C. Fraser, E. A. Wall, S. R. Quake, and M. I. Simon (2010). “Variability in G-protein-coupled signaling studied with microfluidic devices.” *Biophysical Journal* 99.8 (Oct. 2010), pp. 2414–2422.
- Bardwell, L. (2004). “A walk-through of the yeast mating pheromone response pathway”. *Peptides* 25.9 (Sept. 2004), pp. 1465–1476.
- Bashor, C. J., N. C. Helman, S. Yan, and W. A. Lim (2008a). “Using Engineered Scaffold Interactions to Reshape MAP Kinase Pathway Signaling Dynamics”. *Science (New York, NY)*.
- Bashor, C. J., N. C. Helman, S. Yan, and W. A. Lim (2008b). “Using engineered scaffold interactions to reshape MAP kinase pathway signaling dynamics”. *Science (New York, NY)* 319.5869 (Mar. 2008), pp. 1539–1543.
- Burov, S., S. M. A. Tabei, T. Huynh, M. P. Murrell, L. H. Philipson, S. A. Rice, M. L. Gardel, N. F. Scherer, and A. R. Dinner (2013). “Distribution of directional change as a signature of complex dynamics.” *Proceedings of the National Academy of Sciences of the United States of America* 110.49 (Dec. 2013), pp. 19689–19694.
- Cai, L., N. Friedman, and X. S. Xie (2006). “Stochastic protein expression in individual cells at the single molecule level.” *Nature* 440.7082 (Mar. 2006), pp. 358–362.
- Chen, J.-Y., J.-R. Lin, K. A. Cimprich, and T. Meyer (2012). “A Two-Dimensional ERK-AKT Signaling Code for an NGF-Triggered Cell-Fate Decision”. *Molecular Cell* 45.2 (Jan. 2012), pp. 196–209.
- Cheong, R., A. Rhee, C. J. Wang, I. Nemenman, and A. Levchenko (2011). “Information transduction capacity of noisy biochemical signaling networks.” *Science (New York, NY)* 334.6054 (Oct. 2011), pp. 354–358.
- Cohen-Saidon, C., A. A. Cohen, A. Sigal, Y. Liron, and U. Alon (2009). “Dynamics and variability of ERK2 response to EGF in individual living cells.” *Molecular Cell* 36.5 (Dec. 2009), pp. 885–893.

- Coppey, M., A. N. Boettiger, A. M. Berezhkovskii, and S. Y. Shvartsman (2008). “Nuclear trapping shapes the terminal gradient in the *Drosophila* embryo.” *Current biology : CB* 18.12 (June 2008), pp. 915–919.
- Debnath, J., S. K. Muthuswamy, and J. S. Brugge (2003). “Morphogenesis and oncogenesis of MCF-10A mammary epithelial acini grown in three-dimensional basement membrane cultures.” *Methods (San Diego, Calif.)* 30.3 (July 2003), pp. 256–268.
- Fallahi-Sichani, M., S. Honarnejad, L. M. Heiser, J. W. Gray, and P. K. Sorger (2013). “Metrics other than potency reveal systematic variation in responses to cancer drugs”. *Nature Chemical Biology* 9.11 (Nov. 2013), pp. 708–714.
- Feinerman, O., J. Veiga, J. R. Dorfman, R. N. Germain, and G. Altan-Bonnet (2008). “Variability and robustness in T cell activation from regulated heterogeneity in protein levels.” *Science (New York, NY)* 321.5892 (Aug. 2008), pp. 1081–1084.
- Flusberg, D. A., J. Roux, S. L. Spencer, and P. K. Sorger (2013). “Cells surviving fractional killing by TRAIL exhibit transient but sustainable resistance and inflammatory phenotypes.” *Molecular Biology of the Cell* 24.14 (July 2013), pp. 2186–2200.
- Friedman, N., L. Cai, and X. S. Xie (2006). “Linking stochastic dynamics to population distribution: an analytical framework of gene expression.” *Physical review letters* 97.16 (Oct. 2006), p. 168302.
- Janetopoulos, C. and R. A. Firtel (2008). “Directional sensing during chemotaxis.” *FEBS letters* 582.14 (June 2008), pp. 2075–2085.
- Kaern, M., T. C. Elston, W. J. Blake, and J. J. Collins (2005). “Stochasticity in gene expression: from theories to phenotypes.” *Nature Reviews Genetics* 6.6 (June 2005), pp. 451–464.
- Kepler, T. B. and T. C. Elston (2001). “Stochasticity in transcriptional regulation: origins, consequences, and mathematical representations.” *Biophysical Journal* 81.6 (Dec. 2001), pp. 3116–3136.
- Kraskov, A., H. Stögbauer, and P. Grassberger (2004). “Estimating mutual information.” *Physical review. E, Statistical, nonlinear, and soft matter physics* 69.6 Pt 2 (June 2004), p. 066138.
- Lee, R. E. C., S. R. Walker, K. Savery, D. A. Frank, and S. Gaudet (2014). “Fold change of nuclear NF- κ B determines TNF-induced transcription in single cells.” *Molecular Cell* 53.6 (Mar. 2014), pp. 867–879.
- Mehta, P., S. Goyal, T. Long, B. L. Bassler, and N. S. Wingreen (2009). “Information processing and signal integration in bacterial quorum sensing.” *Molecular Systems Biology* 5, p. 325.
- Miller, M., M. Hafner, E. Sontag, N. Davidsohn, S. Subramanian, P. E. M. Purnick, D. Lauffenburger, and R. Weiss (2012). “Modular design of artificial tissue homeostasis: robust control through synthetic cellular heterogeneity.” *PLoS Computational Biology* 8.7, e1002579.
- Nemenman, I., W. Bialek, and R. de Ruyter van Steveninck (2004). “Entropy and information in neural spike trains: progress on the sampling problem.” *Physical review. E, Statistical, nonlinear, and soft matter physics* 69.5 Pt 2 (May 2004), p. 056111.

- Paninski, L. (2006). “Estimation of Entropy and Mutual Information”. *dx.doi.org* 15.6 (Mar. 2006), pp. 1191–1253.
- R: A Language and Environment for Statistical Computing (2010). Vienna, Austria: The R Foundation for Statistical Computing.
- Rowland, M. A., W. Fontana, and E. J. Deeds (2012). “Crosstalk and Competition in Signaling Networks”. *Biophysical Journal* 103.11 (Dec. 2012), pp. 2389–2398.
- Sacan, A., H. Ferhatosmanoglu, and H. Coskun (2008). “CellTrack: an open-source software for cell tracking and motility analysis.” *Bioinformatics (Oxford, England)* 24.14 (July 2008), pp. 1647–1649.
- Sebolt-Leopold, J. S. and R. Herrera (2004). “Targeting the mitogen-activated protein kinase cascade to treat cancer”. *Nature reviews Cancer* 4.12 (Dec. 2004), pp. 937–947.
- Selimkhanov, J., B. Taylor, J. Yao, A. Pilko, J. Albeck, A. Hoffmann, L. Tsimring, and R. Wollman (2014). “Accurate information transmission through dynamic biochemical signaling networks”. *Science (New York, NY)* 346.6215 (Dec. 2014), pp. 1370–1373.
- Shahrezaei, V. and P. S. Swain (2008). “The stochastic nature of biochemical networks.” *Current Opinion in Biotechnology* 19.4 (Aug. 2008), pp. 369–374.
- Shalek, A. K. et al. (2014). “Single-cell RNA-seq reveals dynamic paracrine control of cellular variation.” *Nature* 510.7505 (June 2014), pp. 363–369.
- Shannon, C. E. (1948). “A Mathematical Theory of Communication”. *Bell System Technical Journal* 27 (July 1948), pp. 379–423.
- Spencer, S. L., S. Gaudet, J. G. Albeck, J. M. Burke, and P. K. Sorger (2009). “Non-genetic origins of cell-to-cell variability in TRAIL-induced apoptosis.” *Nature* 459.7245 (May 2009), pp. 428–432.
- Steuer, R., J. Kurths, C. O. Daub, J. Weise, and J. Selbig (2002). “The mutual information: detecting and evaluating dependencies between variables.” *Bioinformatics (Oxford, England)* 18 Suppl 2, S231–40.
- Suderman, R. and E. J. Deeds (2013). “Machines vs. Ensembles: Effective MAPK Signaling through Heterogeneous Sets of Protein Complexes”. *PLoS Computational Biology* 9.10 (Oct. 2013), e1003278.

The Noise is the Signal: Information Flow in Single Cells and Cellular Populations

Supporting Information

Ryan Suderman^{1,*}, John A. Bachman^{2,*}, Adam Smith¹, Peter K. Sorger² and Eric J. Deeds^{1,3}

¹Center for Computational Biology, The University of Kansas, 2030 Becker Dr., Lawrence, KS 66047

²Department of Systems Biology, Harvard Medical School, 200 Longwood Ave., Boston, MA 02115

³Department of Molecular Biosciences, The University of Kansas, 1200 Sunnyside Ave., Lawrence, KS 66047

*These authors contributed equally to this work

Email: Peter K. Sorger - peter_sorger@hms.harvard.edu, Eric J. Deeds - deeds@ku.edu;

Contents

1	Information Theory Calculations	3
1.1	Mutual Information	3
1.1.1	Calculating the mutual information	3
1.1.2	Removing bias due to finite sample size	4
1.1.3	Finding the optimal number of bins	5
1.2	Channel Capacity	8
1.2.1	Unimodal signal distributions	8
1.2.2	Bimodal signal distributions	9
1.2.3	Piece-wise signal distributions	9
1.2.4	Weighting the data	10
2	Additional Experimental Calculations	10
2.1	Control calculations	10
2.2	Population size dependence of single-cell channel capacity	11
2.3	Dose-dependent scaling	11
2.4	Resampling experimental data	11
2.5	Other channel capacities	12

3	Spatial Channel Capacity	12
3.1	Neutrophil motion	14
3.2	<i>Dictyostelium</i> motion	15
4	Simple Model	16
4.1	Choosing signal values	16
4.2	Varying n	17
4.3	Channel capacity saturation with population size	17
4.4	Maximal fractional response	18
4.5	Population channel capacity dependence on signal spacing	20
	References	24

1 Information Theory Calculations

1.1 Mutual Information

The mutual information between two random variables representing a signal, S , and a response, R , is defined as:

$$I(S; R) = \int_S \int_R p(s, r) \log \frac{p(s, r)}{p(s)p(r)} ds dr,$$

where S is a random variable representing the input signal, R a random variable representing the response, $p(s, r)$ is the joint probability distribution for some combination of s and r values, and $p(s)$ and $p(r)$ are the corresponding marginal distributions (1). One of the major difficulties in calculating this quantity from experimental data is the fact that the continuous probability density functions defined above must be estimated on the basis of an inherently discrete data set. As a result, a number of approaches have been developed to obtain unbiased estimates of the mutual information with varying degrees of accuracy (2).

In order to facilitate comparison with earlier results, we employed the same strategy used by Cheong *et al.* (3). This strategy has two main components. First, one defines a set number of “bins” in both the signal values s and response values r . In cases where one is measuring the molecular response of individual cells to a given signal (e.g. nuclear localization of NF- κ B upon treatment with TNF- α , (3)), there are a small number of ligand concentrations used to treat the cells, resulting in a natural discretization of the S variable and a total of S_B bins of signal values. One defines a number of bins for the response (R_B), and uses these bins to estimate the probability of observing some response bin given some signal bin (i.e. $p(r|s)$). A linear extrapolation procedure is then used to estimate the mutual information one would obtain if there were an infinite amount of data in the data set. This extrapolation procedure is described in greater detail in section 1.1.2 below.

One issue with this approach, however, is that the number of bins into which the signal and response values should be divided is not well-defined; using a larger number of bins generally increases the estimated amount of information (3). To combat the potential for overestimation of the mutual information, the second phase of the procedure involves varying the total number of bins in the response variable (and, when appropriate, in the signal value as well) and estimating I for both the experimental data and a set of randomized replicates of the data. This allows one to choose a bin size that maximizes I for the real data while still estimating 0 information for the randomized versions. This element of the procedure is detailed in section 1.1.3. Estimates of the mutual information based on this approach can subsequently be used to calculate the channel capacity by finding the input distribution that maximizes I (1) (section 1.2).

1.1.1 Calculating the mutual information

To calculate the mutual information from our finite data sets, we first created a “contingency table” K based on the data: the rows of this matrix represent the various signal bins, and the columns are the various response bins. Each entry in the matrix is the number of observations from the data that correspond to that particular signal-response pair. The contingency table for a particular experiment might look something like this:

$$K = \begin{matrix} & \begin{matrix} r_1 & r_2 & r_3 & r_4 & r_5 \end{matrix} \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \\ s_4 \end{matrix} & \begin{pmatrix} 6 & 1 & 0 & 0 & 0 \\ 0 & 3 & 4 & 0 & 0 \\ 0 & 1 & 2 & 4 & 0 \\ 0 & 0 & 0 & 2 & 5 \end{pmatrix} \end{matrix}$$

Note that the above table is meant only as an example, and does not contain actual data. One can use the contingency table to calculate the mutual information in terms of the marginal and conditional entropies:

$$I(S; R) = H(S) - H(S|R)$$

$$I(S; R) = - \sum_i^{S_B} p(s_i) \log p(s_i) - \sum_i^{S_B} \sum_j^{R_B} p(s_i, r_j) \log \frac{p(r_j)}{p(s_i, r_j)}$$

where i ranges over the signal bins and j over the response bins in the contingency table (recall that S_B and R_B are the total number of signal and response bins, respectively). Since each entry in the contingency table can be naturally considered a conditional probability, it is helpful to rewrite this equation as:

$$I(S; R) = - \sum_i^{S_B} p(s_i) \log p(s_i) + \sum_j^{R_B} p(r_j) \sum_i^{S_B} p(r_j|s_i) \log p(r_j|s_i).$$

We can then calculate the frequencies from the contingency table entries and substitute these values into the equation. We define N_T as the sum over all entries in the table (i.e. the total number of observations). Since each entry of the matrix, k_{ij} , is the number of instances of signal i that resulted in response j , we can define the total number of observations corresponding to a given signal bin i as $k_{s,i} \equiv \sum_j^{R_B} k_{ij}$. Similarly, we can define the total number of times any particular response bin j was observed as $k_{r,j} \equiv \sum_i^{S_B} k_{ij}$. Given these definitions, we can calculate the mutual information using the following equation:

$$I(S; R) = - \sum_i^{S_B} \frac{k_{s,i}}{N_T} \log \frac{k_{s,i}}{N_T} + \sum_j^{R_B} \frac{k_{r,j}}{N_T} \sum_i^{S_B} \frac{k_{ij}}{N_T} \log \frac{k_{ij}}{N_T}. \quad (1)$$

Equation 1 is used whenever a particular value of I is calculated in the estimation procedure described below (2; 3).

1.1.2 Removing bias due to finite sample size

Although it is straightforward to use equation 1 to calculate the mutual information, the fact that there are a finite number of data points in the contingency table (N_T) can introduce biases into the calculation. To estimate this bias, one can create a smaller data set with N'_T points ($N'_T < N_T$) and calculate I . As has been observed previously (3), as N'_T decreases, bootstrap replicates of the data generate higher values of I . This results in a roughly linear decrease in I as the inverse sample size ($1/N'_T$) decreases (Figure S1).

To correct for this bias, we used the linear extrapolation procedure employed in Cheong *et al.* and other previous studies (3; 4). We chose a range of N'_T values from 80 to 97.5% of the original data, in increments of 2.5%. For each value of N'_T we generated 20 bootstrap replicates, sampling a total of N'_T observations with equal probability from the original data set. We used this randomly subsampled data to generate a new contingency table, and using equation 1 we calculated the average of I across the 20 replicates. We then performed a linear regression of the I vs. $1/N'_T$ relationship (*e.g.* the straight lines in Figure S1). The y-intercept of these lines represents the extrapolation to an infinite data set (*i.e.* $N_T \rightarrow \infty$ implies $1/N_T \rightarrow 0$). All the channel capacities calculated in this work (*e.g.* those reported in Table 1 of the main text) were obtained from these y-intercepts. The errors reported for these values in Table 1, and the error bars in all figures, represent 95% confidence intervals on the linear model's intercept estimate.




Figure S1: Representative linear models for estimating mutual information at infinite sample size with various numbers of response bins. Here we use experimental data from 100 randomly sampled subpopulations of cells where each subpopulation is composed of 1000 cells per each of 19 TRAIL concentrations. In this case we are estimating the mutual information for the fractional quantity of apoptotic cells in response to a uniformly distributed set of TRAIL concentrations. We then calculate the mean mutual information as a function of inverse sample size by taking n independent subsets of the data per sample size. Shown here are the mean values with their corresponding standard error ($n = 20$). Calculation of the linear model's intercept and its standard error provide us with an estimate of mutual information at infinite sample size for a particular number of response bins. The data was resampled using lower percentages (60% to 95% in increments of 5%) than was mentioned in the text for visualization purposes.

1.1.3 Finding the optimal number of bins

Generating the contingency table relies on a particular discretization or binning of the data. As mentioned above, the signal values used to generate experimental data often represent a natural set of signal bins (*e.g.* Figure 1C and D of the main text). The number of response bins to

generate, however, is not clear *a priori*, and the value of R_B has a large impact on estimates of I . On one extreme, if we set $R_B = 1$, all of the signals will give the same responses, resulting in a mutual information of 0. Alternatively, we could choose a number of bins, R_B , so large that every response bin contains exactly one response value in the contingency table:

$$K' = \begin{matrix} & r_1 & r_2 & r_3 & r_4 & r_5 & r_6 & r_7 & r_8 & r_9 & r_{10} & r_{11} & r_{12} \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \\ s_4 \end{matrix} & \begin{pmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 1 \end{pmatrix} \end{matrix}$$

(where again we have used an arbitrary data set as an example). This results in a (spuriously) high mutual information—note that, in this case, if we randomly shuffle the signal value that gives any particular output, we will get the same mutual information:

$$K'_{rand} = \begin{matrix} & r_1 & r_2 & r_3 & r_4 & r_5 & r_6 & r_7 & r_8 & r_9 & r_{10} & r_{11} & r_{12} \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \\ s_4 \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \end{pmatrix} \end{matrix}$$

Since I generally increases with an increasing R_B (note the increasing intercept for the data in Figure S1), we must find an optimal value of R_B that accurately represents the mutual information in the underlying data without artificially inflating it.

Our approach to solving this problem is broadly inspired by previous approaches, particularly that of Cheong *et al.*, with some slight modifications (3; 5). Given some data set, we considered a range of R_B values; this range was independently determined for each data set (see below). For any given R_B value, we generated the bins themselves (i.e. the actual range of response values in the data that belongs to each bin) so that the total number of observations $k_{r,j}$ for each response bin is (roughly) equal across all the bins under a uniform signal distribution (3; 5). We then generated the contingency table and estimated I using the linear extrapolation procedure explained above in section 1.1.2.

Plotting I *vs.* the total number of response bins (Figure S2) does indeed demonstrate that mutual information increases essentially monotonically with increasing R_B . For each value of R_B , we also generated a randomized data set where we shuffled the signal values. To do this, we chose two points in the data set at random, say (s_i, r_j) and (s_k, r_l) and swapped their signal values, generating (s_k, r_j) and (s_i, r_l) ; we repeated this shuffling procedure until the data was completely randomized (3; 5). We generated 10 random replicates of this shuffled data and calculated I for each one of them; the value of I in these randomized data sets also increases with increasing R_B , eventually generating significantly non-zero mutual information where there should be none (Figure S2).

Cheong *et al.* obtained an optimal range of bin numbers for each data set *via* visual inspection of plots like those in Figure S2 (3). While this is an effective approach, the large number of data sets




Figure S2: Here we show a representative graph of mutual information (I) as a function of the number of response bins. The data is the same from which Figure S1 is generated and each point shown here is the y -intercept and its corresponding 95% confidence interval retrieved from the linear extrapolation procedure outlined in section 1.1.2. In red is one of ten randomizations of the actual data set, shown in black. If we consider only this randomization, the maximum number of bins that is eligible under our criteria (*i.e.* estimated mean I - its 95% confidence interval ≤ 0 for the randomized data) is 56 which produces $I = 3.40 \pm 0.01$

and variants in our case prevented us from visually analyzing every case. We thus defined a uniform criterion for choosing the optimal number of bins, defined as the value of R_B that gives the largest value of I , subject to the constraint that the 95% confidence interval from the corresponding randomized data *must include* 0. In other words, we chose an R_B that maximizes the I in the data, but where the randomized data gives mutual information that is not significantly greater than 0.

The range of R_B values that provides this maximum depends on the total number of data points (N_T) and on the amount of information present in the data itself; it is thus difficult to define a uniform range of bin numbers to consider for every data set. To find a useful range of bin numbers, we considered a range of bin numbers (say, 40 to 120). If the minimum bin number in that set gave non-0 information for the randomized data, we extended the range of bin numbers to smaller values. If the maximum bin number in that range had 0 information in the randomized data sets, we expanded the range to larger numbers. The vast majority of cases we considered had optimal values of R_B below 100.

The discussion above assumes that S_B is fixed at a (relatively) small number of signal values used to generate experimental data. In some of the systems we considered, however, we needed to find an optimal set of signal bins in addition to response bins. This was particularly true of spatial quantities like the angle between the bacterium and the neutrophil (see section 3). In those cases, we also defined a range of S_B values, and tried all relevant (S_B, R_B) pairs to find a set of bin

numbers that maximized I , subject again to the constraint that the randomized data gave I values that were not significantly greater than 0.

1.2 Channel Capacity

As mentioned in the main text, the channel capacity is the supremum of the mutual information over all possible signal distributions:

$$C = \sup_{p_S(s)} I(S; R).$$

Estimating the channel capacity thus involves using the estimate of mutual information obtained from the procedure defined in section 1.1 to search the space of signal distributions and find the one that maximizes I . Since the set of such distributions is obviously infinite, an exhaustive search of all well-defined signal distributions is impossible. Following the example of Cheong *et al.* (3), we implemented a grid-based search, limited to a set of unimodal and bimodal Gaussian distributions in addition to the initial uniform distribution of signal values. We also added a series of piece-wise functions designed to uniformly weight a subset of the total sampled signal values. The details of this procedure are described below.

1.2.1 Unimodal signal distributions

We generated a range of unimodal signal distributions of the form:

$$G_U(s) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(s-\mu)^2}{2\sigma^2}}.$$

Since we can sample only a subset of possible signal distributions we limit the potential mean values, μ , to a set of 4 evenly spaced values between the minimum and maximum signal values in the data set (S_{min} and S_{max} , respectively):

$$\mu \in \{f \cdot (S_{max} - S_{min}) + S_{min} \mid f \in \{0.2, 0.4, 0.6, 0.8\}\}.$$

For each μ we generate a range of σ values by calculating a maximum standard deviation:

$$\sigma_{max} = \min\left(\frac{\mu - S_{min}}{3}, \frac{S_{max} - \mu}{3}\right).$$

This constrains the signal distributions so that at least 99% of the area under the distribution falls between S_{min} and S_{max} and allows us to use a range of standard deviation values by sampling increasing fractions of σ_{max} :

$$\sigma \in \{f \cdot \sigma_{max} \mid f \in \{0.2, 0.4, 0.6, 0.8, 1\}\}.$$

1.2.2 Bimodal signal distributions

We also implemented a range of bimodal signal distributions of the form:

$$G_B(s) = \frac{w_0}{\sigma_0\sqrt{2\pi}} e^{-\frac{(s-\mu_0)^2}{2\sigma_0^2}} + \frac{w_1}{\sigma_1\sqrt{2\pi}} e^{-\frac{(s-\mu_1)^2}{2\sigma_1^2}}$$

where w_0 and w_1 are weighting coefficients such that $w_0 \in \{0.4, 0.5, 0.6\}$ and $w_1 = 1 - w_0$. In order to construct these distributions we first defined a minimum difference between μ_0 and μ_1 :

$$\mu_D = \frac{S_{max} - S_{min}}{5}.$$

We used μ_D to construct a series of pairs (μ_0, μ_1) such that

$$\mu_D + S_{min} \leq \mu_0 < S_{max}$$

$$\mu_0 + \mu_D \leq \mu_1 < S_{max}$$

and μ_0 is incremented in steps of μ_D . Similarly to the unimodal signal distributions, both means μ_0 and μ_1 have multiple, evenly spaced standard deviations, σ_0 and σ_1 that are fractions of some maximum standard deviations, $\sigma_{0,max}$ and $\sigma_{1,max}$. These values are constrained so that these distributions have both a local minimum between μ_0 and μ_1 and 99% of their area between S_{min} and S_{max} :

$$\sigma_{0,max} = \min\left(\frac{\mu_1 - \mu_0}{4}, \frac{\mu_0 - S_{min}}{3}\right)$$

$$\sigma_{1,max} = \min\left(\frac{\mu_1 - \mu_0}{4}, \frac{S_{max} - \mu_1}{3}\right).$$

The individual σ_0 and σ_1 values are then:

$$\sigma_0 = f \cdot \sigma_{0,max}$$

$$\sigma_1 = f \cdot \sigma_{1,max},$$

where $f \in \{0.2, 0.4, 0.6, 0.8, 1\}$.

1.2.3 Piece-wise signal distributions

Given a range of signal values (and thus signal bins), we generated a series of piece-wise probability distributions for all combinations of signal bins s_x and s_y where the values in bin s_x are lower than those in bin s_y (see next section, 1.2.4). We denote this inter-bin relationship with common inequality operators (*i.e.* $s_x < s_y$). The probability of finding a signal value in a bin s_i can then be determined from functions of the form:

$$PW(s) = \begin{cases} 0 & : s_i < s_x \\ 0 & : s_i > s_y \\ \frac{1}{j-i+1} & : s_x \leq s_i \leq s_y \end{cases}$$

1.2.4 Weighting the data

With this set of new unimodal and bimodal signal distributions, we can determine how the mutual information of a particular data set varies with different signal distributions in order to estimate the channel capacity. To do this, we modified the original contingency table (constructed assuming a uniform signal distribution) in order to recalculate the mutual information according to each new signal distribution. Each signal bin s_i corresponds to a range of signal values between, say, $s_{i,min}$ and $s_{i,max}$ and yields a corresponding number of observations in the contingency table, k_{ij} for each response bin r_j . For any new unimodal or bimodal signal distribution $G_A(s)$, we calculated the new value for this entry in the contingency table $k'_{ij} = \frac{p'(s_i)k_{ij}}{p(s_i)}$ where

$$p'(s_i) = \int_{s_{i,min}}^{s_{i,max}} G_A(s) ds$$

is the new probability of observing some signal value s_i and $p(s_i)$ is the original, uniform probability of observing that signal bin s_i . To determine the correct signal probability from our piece-wise weighting scheme we simply calculate

$$p'(s_i) = PW(s_i).$$

We can use this to generate a new contingency table, and calculate the relevant quantities:

$$N'_T = \sum_i^{S_B} \sum_j^{R_B} k'_{ij}, \quad k'_{s,i} = \sum_j^{R_B} k'_{ij}, \quad k'_{r,j} = \sum_i^{S_B} k'_{ij}.$$

The procedure produces a new contingency table that has approximately the same number of entries as the original one, $N'_T \approx N_T$. For each distribution $G_A(s)$ that we considered, we used the procedures described in section 1.1 to estimate the mutual information for that particular distribution. The maximum mutual information over all signal distributions calculated is our estimate of the channel capacity C . As mentioned above, the errors reported for C represent the 95% confidence interval for the intercept estimated by the linear extrapolation procedure (section 1.1.2).

2 Additional Experimental Calculations

2.1 Control calculations

As mentioned in the main text, we examined the channel capacities between the activities of the initiator caspase (IC; cleaved caspase 3) and both the effector caspase (EC; cleaved PARP) and terminal cellular phenotype. We found that the IC to EC channel capacity exceeded 1.2 bits with a IC to cell fate channel capacity of approximately 0.8 bits (Table S1, entries 9 & 10). This confirms the nature of IC as an intermediate component in the TRAIL signaling network due to the relative increase in information when using IC as the input distribution to the channel capacity calculation instead of TRAIL. As an addendum, these calculations also highlight the role that upstream receptor-based signaling components play in generating noise within the network

as enforced by the data processing inequality (6).

2.2 Population size dependence of single-cell channel capacity

Given our large data set, we investigated how channel capacity would vary for individual cells as a function of population size; as mentioned in the main text, the population-level channel capacity has a clear dependence on the size of the population. We expected to find that as the sample size increases the estimators describing the response distribution will be sufficiently accurate to prevent the need to calculate the channel capacity from the entire data set of over 1.2 million cells (which is computationally expensive). We confirmed this empirically, upon calculation of the single-cell channel capacity for increasing subsets of our FACS-generated data, using sample sizes of 500, 1000, 2000, and 4000 cells per TRAIL concentration.

2.3 Dose-dependent scaling

In our data we observed that IC activity levels were substantially higher in dead than live cells, most likely due to the variety of positive feedback mechanisms present in the caspase cascade (7). Because this additional cleavage of caspase 3 in dead cells occurs downstream of the cell's commitment to apoptosis, it could be considered a consequence of the cell's phenotypic outcome rather than as an intermediate factor contributing to it. Since the channel capacity estimation is time-dependent (and capturing the exact moment of cell death for every cell is technologically infeasible), we proceeded to examine the impact that post-commitment IC activity has on our estimates for channel capacity between TRAIL dose and IC activity level. We therefore performed our analysis separately for live and dead cells by partitioning them into these two groups according to the threshold effector caspase response, t_{EC} (8); we calculated this quantity by estimating the minimum density between the two peaks of the bimodal EC activity distribution: $\log_{10}(t_{EC}) = 2.85 \pm 0.05$ (9). We then plotted the dose response data to determine how response varies with TRAIL concentration. These plots show clearly that only initiator caspase activity scales with TRAIL dose, and it does so only among living cells (Figure S3). The channel capacity between TRAIL and IC activity in living cells is approximately 0.67 bits (Table S1), nearly the same as the channel capacity between TRAIL and IC activity in all cells. Mean effector caspase activity in living and dead cells in addition to mean initiator caspase activity in dead cells does not significantly vary for differing doses of TRAIL (Figures S3 and S4) and as a result, we did not calculate the channel capacity for these dose-response relationships.

2.4 Resampling experimental data

In order to calculate the channel capacity for the pheromone signaling network in yeast, we reconstructed dose-response data shown by Bashor *et al.* in their Figure 4D (10). In this case the signal distribution was a set of logarithmically spaced α -factor concentrations and the corresponding response was *pFUS1*-GFP fluorescence. With this data we constructed a series of dose-dependent Gaussian distributions defined by the mean and standard deviation of the *pFUS1*-GFP response given some α -factor concentration. From these distributions we sampled




Figure S3: Initiator caspase activity scales with TRAIL among living cells. Shown for each TRAIL concentration are the sample mean and standard deviation ($n \approx 60,000$ cells)

100 values for each of 10 α -factor concentrations in order to construct a dose-response data set from which we could estimate the channel capacity (Figure S5). We similarly performed this procedure for calculating the population-level channel capacity for the set of MCF10A cells shown in Figure 3B of the main text. In this case, the mean and standard deviation for a particular TRAIL dose refer to the number of living cells in a given population.

2.5 Other channel capacities

Table S1 lists all channel capacity calculations for experimental data that are relevant to this work; it contains the values reported in Table 1 in the main text in addition to other calculations performed in this work and others. The entries containing population-level channel capacities were calculated using 100 independently sampled populations.

3 Spatial Channel Capacity

In order to calculate the spatial channel capacity between a motile cell undergoing chemotaxis and its target that is producing some chemical gradient, we constructed signal/response pairs from angles between the cell and its target. We used the CellTrack program developed by Sacan *et al.* (15) to output text files containing frame-by-frame coordinates for the edges of the cells and their centers of mass (COM). We used these coordinates to calculate time-dependent signal and response angles.

Signal (molecular)	Response (molecular)	C (bits)	Data	Calculation
1. TNF	NF- κ B	0.92 ± 0.01	(3)	(3)
1.1. TNF	ATF-2	0.85 ± 0.02	(3)	(3)
1.2. TNF	NF- κ B & ATF-2	1.05 ± 0.02	(3)	(3)
2. PDGF	NF- κ B	0.67 ± 0.01	(3)	(3)
2.1. PDGF	ATF-2	0.74 ± 0.01	(3)	(3)
2.2. PDGF	NF- κ B & ATF-2	0.81 ± 0.02	(3)	(3)
3. EGF	Erk (fold-change)	0.60 ± 0.03	(11)	(3)
4. UDP	Peak Ca^{2+}	1.22 ± 0.03	(12)	(3)
4.1. UDP	Integrated Ca^{2+}	1.07 ± 0.02	(12)	(3)
5. TNF	A20 transcripts	0.62 ± 0.08	(13)	this work
6. TRAIL	Casp-8 activity	0.66 ± 0.02	this work	this work
7. TRAIL	Casp-8 activity (live cells)	0.67 ± 0.01	this work	this work
8. TRAIL	Casp-3 activity	0.33 ± 0.01	this work	this work
9. Casp-8 activity	Casp-3 activity	1.23 ± 0.02	this work	this work
10. Casp-8 activity	cell decision	0.82 ± 0.01	this work	this work
11. α -factor	<i>pFUS1</i> -GFP	2.26 ± 0.05	(10)	this work
Signal (position)	Response (molecular)	C (bits)	Data	Calculation
12. Embryo perimeter	Phosphorylated Erk	1.61 ± 0.05	(14)	(3)
Signal (position)	Response (motion)	C (bits)	Data	Calculation
13. Bacterium	neutrophil motion	2.14 ± 0.22		this work
14. cAMP	<i>Dictyostelium</i> motion	2.24 ± 0.07	Firtel Lab	this work
Signal (molecular)	Response (population)	C (bits)	Data	Calculation
15. TRAIL	% dead (HeLa cells)	3.46 ± 0.03	this work	this work
16. TRAIL	% dead (MCF10A cells)	3.36 ± 0.03	this work	this work

Table S1: Channel capacities for experimental data sets




Figure S4: Effector caspase activity is invariant with TRAIL among both living and dead cells. The data set is identical to that in Figure S3.

3.1 Neutrophil motion

We initially analyzed the motion of a neutrophil that is “chasing” a bacterium from a classic movie taken in the 1950s (see supplemental files, or [this website](#)). For the purposes of our calculation, we assume that the neutrophil is in fact following a chemical gradient generated by the bacterium. In this case, the signal corresponds to the angle, termed θ_1 , between the bacterium at a particular frame in the movie, time t , and the neutrophil at another time $t + \Delta t_1$. The subsequent response angle, θ_2 , is that of neutrophil motion between time $t + \Delta t_1$ and time $t + \Delta t_1 + \Delta t_2$. These angles then comprise the signal and response distributions used to calculate the channel capacity of the system. A visual representation of this calculation can be seen in the main text (Figure 4).

We calculated the signal and response angles between the neutrophil COM and bacterial COM relative to the x -axis unit vector in the Cartesian coordinate system. This method is similar to one outlined by Burov *et al.* derived to provide more directional information than mean squared displacement for analysis of random walks (16). We employ a “windowed” data collection method; given some starting time, t , we calculate an arbitrary signal and response angle pair, requiring information from time points, $t + \Delta t_1$ and $t + \Delta t_1 + \Delta t_2$. In our windowed data collection, the next pair of angles is calculated using t incremented by one frame: $t = t + 1$. To confirm that the calculated channel capacity was not an artifact of the chosen time delay values, Δt_1 and Δt_2 , we explored the nearby $(\Delta t_1, \Delta t_2)$ -space and discovered that the channel capacity is relatively robust to Δt_1 and Δt_2 as seen in Figure S6.

We also investigated other potential methods for calculating the signal angles. Our first alternative calculation involved introducing a *radius of detection*, R , about the neutrophil COM. The bacterial COM that is nearest, but outside, the circle defined by the neutrophil COM and R , **regardless of the time the bacterium occupies that position**, is then chosen for the signal angle calculation. This eliminates the need for Δt_1 . Our second alternative calculation involved




Figure S5: Resampled data from Figure 4D in Bashor *et al.* (10).

Calculation	Parameters	Channel Capacity
Initial	$\Delta t_1 = 3, \Delta t_2 = 9$	2.14 ± 0.22
Radius of Detection	$R = 70, \Delta t_2 = 10$	2.21 ± 0.27
Leading Edge	$\Delta t_2 = 1$	1.99 ± 0.11

Table S2: Alternate angle calculation methods for spatial channel capacity estimation

determination of the distance between the neutrophil's *leading edge* and the bacterial COM at some time t . The point on the neutrophil's edge resulting in the smallest distance was then used for the signal angle calculation. This was done over all time points, eliminating the need for both Δt_1 and R . In both alternative calculation methods the response angles are calculated from the neutrophil's positions at t and $t + \Delta t_2$, and the estimated channel capacities are seen in Table S2:

3.2 Dictyostelium motion

The next movie we analyzed is that of a *Dictyostelium* cell following a cAMP gradient (see supplemental files or [this website](#)). In this movie, *Dictyostelium* responds to cAMP introduced by a pipette tip which changes location periodically. Since the pipette tip remains stationary between location shifts, we can employ our original calculation used for the neutrophil/bacterium data and omit the Δt_1 parameter (the largest channel capacity occurs when $\Delta t_2 = 15$). This omission is valid since θ_1 is identical for a range of θ_2 values (*i.e.* no motion in the gradient source/change in the signal).




Figure S6: Channel capacity as it depends on Δt_1 and Δt_2 . Sampling time delay values near to those producing the maximum channel capacity ($\Delta t_1 = 16$ and $\Delta t_2 = 3$) results in similar channel capacities.

4 Simple Model

As mentioned in the main text, the initial model takes the form:

$$R = (R_{max} - R_{min}) \cdot \frac{S^n}{S^n + K^n} + R_{min} + \epsilon \quad (2)$$

where the normally-distributed noise term $\epsilon \sim N(0, \sigma)$ depends on some chosen standard deviation, σ . The parameter values chosen for the base model (shown in Figure 3) are as described in the Materials and Methods section of the main text: $K = 10$, $n = 6$, $R_{max} = 30$, and $R_{min} = 20$. For all models discussed in the paper, the response threshold governing an individual cell's fate is positioned such that half of the signal values produce mean responses below the threshold and half produce mean responses above the threshold.

4.1 Choosing signal values

We selected evenly-spaced signal values to achieve responses 10% above the minimum response and 10% below the maximum response (the *transition zone*):

$$0.1 \cdot (R_{max} - R_{min}) + R_{min} \leq R \leq 0.9 \cdot (R_{max} - R_{min}) + R_{min}$$

Through simple algebra we show that the resulting minimum and maximum signal values, (S_{min} and S_{max} , respectively) are:

$$S_{min} = K \cdot \sqrt[n]{\frac{1}{9}}$$




Figure S7: Single-cell channel capacity with respect to noise for a range of n values. The starred $n = 6$ denotes the value used in all other calculations based on this model. There is minimal difference between models where $n > 2$ and even the model with $n = 2$ displays qualitatively similar behavior to the others. Error bars denote 95% confidence about the intercept estimate (see section 1.1.2)

and

$$S_{max} = K \cdot \sqrt[n]{9}$$

This prevents selection of signal values that would produce extremely high or low responses, since sampling more of these responses relative to intermediate responses would reduce the channel capacity.

4.2 Varying n

In order to determine the effect of our chosen $n = 6$ on this model's channel capacity (both single-cell and population-level), we varied n between 2 and 10. We see in general from Figures S7 and S8 that this variation produces minimal difference between models; qualitatively, models with different n are nearly identical.

4.3 Channel capacity saturation with population size

As discussed in the main text, the fraction of a group of cells making a particular signal-dependent decision is the statistic used to determine collective response for the population-level calculation of mutual information. By calculating this statistic over a number of replications, we effectively construct a sampling distribution for the fractional response to some arbitrary signal. Since the standard deviation of this distribution (*i.e.* the standard error of the statistic) is dependent on sample size, we observe an inverse correlation between the size of the




Figure S8: Population-level channel capacity with respect to noise for a range of n values. Again we see little difference between models with different n with the minor exception of $n = 2$. We do observe a slight shift in the amount of noise producing maximal channel capacity, but the qualitative trends are essentially identical. Error bars are as in Figure S7.

population and the standard error of the fractional response (Figure S9). If we restrict our data set to those values in the increasing regime of the dose-response curve (the *transition zone*, see Section 4.1), increasing the population size results in the channel capacity approaching its theoretical maximum of $C = -\log_2(\frac{1}{N})$ bits (*i.e.* the entropy of the signal distribution in the transition zone) where N is the number of signal values in the transition zone. We observe this channel capacity saturation in Figure 3D of the main text.

4.4 Maximal fractional response

At high levels of noise, we observe another interesting feature of the population response: the inability to effect a universal population response at arbitrarily high signal levels. In other words, no matter how much signal is present in the environment, there will still be a fraction of cells in a population that does not respond. This is plainly observed graphically in Figure S10; even as the response saturates at low and high signal levels, there is sufficient noise such that a subpopulation of cells at these signal levels fall above and below the threshold, respectively. The corresponding population dose-response curve thus exhibits saturating, incomplete responses both at low and high signal levels (Figure S11).




Figure S9: Population-level dose-response curves for multiple population sizes




Figure S10: Single cell response curve at high noise values




Figure S11: Population response curve at high noise values (population size = 1000)

4.5 Population channel capacity dependence on signal spacing

As mentioned previously, we restricted sampling signal space (on the individual cell level) to the region generating responses between 10% and 90% of the model's maximum response, since the majority of the information resides in this section of the single-cell dose-response curve (see Section 4.1). However, as the noise decreases on the single-cell level, the shape of the population dose-response curve changes, becoming more switch-like and ultimately shrinking the signal range across which the population-level transition occurs (Figure S12). If we then engage in a similar strategy for the population dose-response curve by sampling a fixed number of signal values corresponding to responses between 10% and 90% of the maximal *population* response for a given noise value, we can keep the population-level channel capacity constant as noise approaches 0, as shown in Figure S12. It thus appears that generating high population-level channel capacity requires only an exceedingly small level of noise, so long as the number of signal values sampled remains constant across the (potentially very narrow) region over which the majority of the transition occurs. A similar numerical experiment reveals that by simply increasing the number of signal values sampled in the original signal range (defined by the signal values corresponding to the single-cell transition region), the amount of noise required to reach the maximal population-level channel capacity decreases (Figure S13). It is thus unclear why cells might have evolved high levels of noise to control population-level responses, when these results suggest that any non-zero level of heterogeneity would suffice.

It is important to note, however, that maintaining a high population-level channel capacity for arbitrarily low levels of noise requires increasingly smaller spacing between individual signal values (see Figure S12). This is problematic for two key reasons:

1. The signal (defined as a quantity of some extracellular signaling factor at a particular point




Figure S12: By sampling evenly-spaced signal values in the population dose-response transition zone (which varies given some level of noise in the individual cell), we observe that the population-level channel capacity can be maintained at a constant value as noise approaches 0.

in time) is inherently discrete since these stimuli are ultimately composed of countable molecules. As such, there is some (potentially very small) level of noise below which one cannot maintain high population-level channel capacities, simply because one cannot achieve a spacing between signal values smaller than adding one molecule of signal at a time.

2. The existence of variability in the signal itself interferes with very small signal spacing in physically realistic systems.

Regarding the (perhaps more biologically relevant) second issue, consider an experimentalist exposing multiple populations of cells to some stimulus in order to measure the fractional response of the population that makes some decision. In this example, there will inevitably be some variation in the quantity of signaling factor to which these independent populations are exposed due to experimental error. Alternatively, if we consider an *in vivo* scenario in which hormones or cytokines are distributed to the various tissues of an organism, we know that these molecules are themselves produced by other cells. Since cytokine production is a stochastic process, the amount of signal to which specific cells within a tissue are exposed will be a variable quantity.

To confirm that the existence of variability in the signal does in fact produce optimal population-level channel capacity at non-trivial levels of single-cell noise, we introduced another noise term (the *signal* noise, ϵ_s , as opposed to the original *response* noise, ϵ) governing the limit of signal accuracy for our populations of simulated cells. As an example, consider signal noise equal to 1% of the signal value that produces an average response corresponding to the decision-making




Figure S13: Increasing the number of sampled signal values in the transition zone of the *individual* cell's dose-response curve (*i.e.* independently of the population dose-response transition zone) results in a decrease in optimal noise level and an increase in channel capacity which appears to approach some limit.

threshold. This alternate form of the model (modified from Equation 2) has the following form:

$$R = (R_{max} - R_{min}) \cdot \frac{(S + \epsilon_s)^n}{(S + \epsilon_s)^n + K^n} + R_{min} + \epsilon$$

where the signal noise term ϵ_s is normally distributed and is sampled independently for each *population* of cells: $\epsilon_s \sim N(0, \sigma)$. This procedure simulates the previously discussed example of the experimentalist's exposure of signal to multiple cellular populations, and it is distinct from the application of the original noise term, ϵ , that was applied to each individual cell. We then calculated the population-level channel capacity for data sets in which the *population-level* transition zone is fixed, by altering the signal space density so that the responses corresponding to these signals fall between 10% and 90% of the maximal population response. Using this data, we characterized the impact of signal detection limits on the population-level channel capacity with respect to different levels of ϵ (Figure S14). As expected, we see that slight error in the signals to which the populations are exposed reduces the population-level channel capacity at low levels of response noise. Thus, although very small levels of single cell noise could theoretically produce high population-level channel capacity, the presence of signal variability in physically realistic systems requires higher levels of noise, so that signal values can be spaced at reasonable distances from one another while maintaining low variability in the population-level response.




Figure S14: Implementing a limit on signal resolution in the form of a fixed level of signal noise (ϵ_s) results in a positive correlation between noise level and channel capacity on the population level. In this figure ϵ_s is one percent of the signal value corresponding to the decision-making response threshold. We also note that this data exhibits a population-level $C < 1$ bit when the single-cell noise is 0, and that this differs from the data in the main text and Figure S13. This results from the presence of relatively high signal noise given the sharp transition region that occurs in the population response with negligible single cell (response) noise (Figure S12) as compared to the lack of signal noise in other mentioned data sets.

References

1. C. E. Shannon, *Bell System Technical Journal* **27**, 379 (1948).
2. L. Paninski, *Neural Computation* **15**, 1191 (2003).
3. R. Cheong, A. Rhee, C. J. Wang, I. Nemenman, A. Levchenko, *Science (New York, NY)* **334**, 354 (2011).
4. S. Strong, R. Koberle, R. de Ruyter van Steveninck, W. Bialek, *Physical Review Letters* **80**, 197 (1998).
5. R. Steuer, J. Kurths, C. O. Daub, J. Weise, J. Selbig, *Bioinformatics (Oxford, England)* **18 Suppl 2**, S231 (2002).
6. T. M. Cover, J. A. Thomas, *Elements of Information Theory* (Wiley-Interscience, New York, NY, USA, 1991).
7. J. G. Albeck, J. M. Burke, S. L. Spencer, D. A. Lauffenburger, P. K. Sorger, *PLoS Biology* **6**, e299 (2008).
8. J. G. Albeck, *et al.*, *Molecular cell* **30**, 11 (2008).
9. *R: A Language and Environment for Statistical Computing* (The R Foundation for Statistical Computing, Vienna, Austria, 2010).
10. C. J. Bashor, N. C. Helman, S. Yan, W. A. Lim, *Science (New York, NY)* **319**, 1539 (2008).
11. C. Cohen-Saidon, A. A. Cohen, A. Sigal, Y. Liron, U. Alon, *Molecular cell* **36**, 885 (2009).
12. X. R. Bao, I. D. C. Fraser, E. A. Wall, S. R. Quake, M. I. Simon, *Biophysical journal* **99**, 2414 (2010).
13. R. E. C. Lee, S. R. Walker, K. Savery, D. A. Frank, S. Gaudet, *Molecular cell* (2014).
14. M. Coppey, A. N. Boettiger, A. M. Berezhkovskii, S. Y. Shvartsman, *Current biology : CB* **18**, 915 (2008).
15. A. Sacan, H. Ferhatosmanoglu, H. Coskun, *Bioinformatics (Oxford, England)* **24**, 1647 (2008).
16. S. Burov, *et al.*, *PNAS* **110**, 19689 (2013).

THIS THESIS WAS TYPESET using \LaTeX , originally developed by Leslie Lamport and based on Donald Knuth's \TeX . The body text is set in 11 point Egenolff-Berner Garamond, a revival of Claude Garamont's humanist typeface. A template that can be used to format a PhD dissertation with this look & feel has been released under the permissive AGPL license, and can be found online at github.com/asm-products/Dissertate or from its lead author, Jordan Suchow, at suchow@post.harvard.edu.