Information and Learning in Mechanism Design

Information and Learning in Mechanism Design

A dissertation presented

by

Jonathan Libgober

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2018

Dissertation Advisors: *Author:*
Professor Eric Maskin **Jonathan Libgober**
Professor Drew Fudenberg

Information and Learning in Mechanism Design

Abstract

This dissertation studies the design of mecanisms in settings where information acquisition or communication are significant features of the environment.

The first chapter, coauthored with Xiaosheng Mu, studies a dynamic pricing model where buyers have the ability to learn about their value for a product over time. A seller commits to a pricing strategy, while buyers arrive exogenously and decide when to make a one-time purchase. The seller seeks to maximize profits against the worst-case information arrival processes. We show that a constant price path delivers the optimal profit, which is also the optimal profit in an environment where buyers cannot delay.

The second chapter develops a model of costly information acquisition, focusing on an application to scientific research. It shows that non-transparency can induce a scientist to undertake a costlier but more informative experiment if it also enables her to commit to acting scrupulously. Using this insight, this chapter demonstrates the general existence of non-degenerate experiment costs such that greater transparency in scientific methodology results in research choices that are worse for those interested in the results.

The third chapter develops a model of delegated project choice with mulitple agents, considering the impact of competition in these settings. Under an alignment assumption (which ensures the optimality of full discretion in the single-agent case), optimal mechanisms entail stochastic agent choice but deterministic project choice. Without alignment, deterministic project choice may be suboptimal. Without the ability to randomize allocations, competition can be harmful for the principal.

Contents

List of Figures

Acknowledgments

While the information acquisition and learning studied in this dissertation is not intended to refer to my own, I am sincerely grateful to those who made that possible.

I am indebted to my primary advisors, Drew Fudenberg and Eric Maskin. Throughout my entire time at Harvard, they have been the two most constant and unwaivering supporters I have had in my academic life. I also owe much to my other committee members, Jerry Green and Ben Golub, for their guidance and support. Ben's support in particular went above and beyond what any student could reasonably expect, and his support was a significant factor in my having made it through the process successfully.

I am particularly grateful to my coauthor, officemate, and friend, Xiaosheng Mu, not just for his joint responsibility for bringing the first chapter into existence, but for his class and friendship throughout graduate school and in particular in the lead up to the job market. Working with him was undoubtedly the highlight of my career at Harvard. A brilliant economist and a wonderful person, I have no doubt he will continue to be a star in the profession.

I also feel special gratitude to my undergraduate mentors, in particular Hugo Sonnenschein and Paul Sally, for having sparked my academic curiosity which has culminated in this dissertation. I will always look up to them as a guide for how to conduct myself as a researcher, and hope that someday I will be able to inspire students as much as they inspired me. I also owe much to Grace Tsiang and Victor Lima, for having cultivated a wonderful undergraduate economics community at Chicago.

In my view, no economics thesis from Harvard would be complete without an acknowledgment of how crucial the support of Brenda Piquet was. Everyone knows how she consistently goes above and beyond for the department, but I also feel very grateful for her support throughout many different phases of graduate school. I also owe special thanks to Kiran Gajwani, who made teaching Ec985 a highlight of my time at Harvard.

Outside of my committee, I was fortunate to have been shaped and influence by other faculty members at Harvard. I deeply appreciate all the guidance I received from Tomasz

Strzalecki on many dimensions, and his support on the first chapter in particular was crucial in helping it take its current shape. Scott Kominers has been one of the longest and most consistent sources of guidance, and conversations with him were helpful for each of my three chapters. The opinions of David Laibson, Matthew Rabin and Shengwu Li also helped shape how I think about theory today. And I appreciate having been able to get to know Claudia Goldin, Ed Glaeser, and Gautam Rao.

I was very lucky to learn from amazing visitors during my time at Harvard as well. Johannes Hörner went above and beyond what was expected of a visitor, and his support played a large role in both the first and second chapter. Stephen Morris' course was one of the best I took at Harvard and formed an early inspiration for the first chapter. Jason Hartline and Michihiro Kandori also shaped the department during their visits and I appreciate that they became involved in the theory community.

The Games and Markets community throughout my entire time at Harvard—Yuhta Ishii, Ryota Iijima, Annie Liang, Mira Frick, Morgan McClellon, Gleb Romanyuk, Jetlir Duraj, Kevin He, Krishna Dasaratha, Shoshana Vasserman, Linh Tô, Neil Thakral, Rob Siliciano, Maria Voronina, Aubrey Clark, Kentaro Tomoeda, Sangram Kadam, Michael Thaler, Marc Kaufmann, Eduard Talamas, Ran Shorrer, Nicola Rosaia and Nir Hak—provided a home where each of my dissertation chapters were presented. I thank them for being part of my life in graduate school.

I was fortunate to be able to complete my studies with a wonderful classmates and scholars. I am especially grateful for the support of my officemates (Xavier Jaravel, Anastassia Fedyk, Kirill Borusyak, Sid George and Ellora Derenoncourt), roommates (Gabriel Unger, Nihar Shah and Aaron Ewall-Wice from MIT Physics), and other close friends who have supported me throughout graduate school, particularly Heather Sarsons, Andy Garin (and of course Allison Bernstein as well), and from MIT Econ, Ben Roth and Vivek Bhattacharya. I was fortunate to have gone through the first year with some wonderful friends—Luca Maini, David Choi, Vitaly Bord, Brian Chen, Yin Chen, Max Eber, Nathan Hipsman, Stephanie Lo, Eduardo Montero, Josh Feng, Alex Segura, Argyris Tsiaras, Edoadro Teso, Peter Tu, Wentao

Xiong, Carl Veller, Kevin Pan, Carmen Wang—and to have gone through the job market with still others—Eben Lazarus, Alonso de Gortari, John Coglianese, Laura Blattner, Jann Spiess, Daniel Lewis, Yueran Ma and Yuxiao Huang.

I especially appreciate the help I received from Guillaume Pouliot and Jeff Picel, who helped advise me both in undergrad and graduate school. I'm very grateful that Harry Pei has been an important part of Drew's student community. Despite neither being at Harvard nor a micro theorist, Paul Ho has also been a constant helpful presence. Eliot Abrams has been as close as family and listened to many of my ideas both before graduate school and during. And the academic community in Cambridge provided a places where I could thrive, and I was fortunate to develop friendships with Ashley Craig, Chenzi Xu, Rachel Cao, Priya Shanmugan, Søren Henn, Dan Sullivan, Sara Lowes, Andrey Fradkin, Mitra Akhtari, Rachael Meager, Will Diamond, Emma Harrington, Zoë Hitzig, John Firth, Jing Li, Frank Schilbach, Fanyin Zheng, Jack Willis, David Baqaee, Natalie Bau, Tzachi Raz, Seth Blumberg, Alex Peysakhovich, Anna Gumpert, Jonathan Roth, Arielle Bernhardt, Natalia Emanuel, Mark Shepard, Pascal Noel, Raissa Fabregas, Giselle Montamat, Elisa Rubbo, Tommaso Denti, Ludwig Straub, Cuicui Chen, Lisa Xu, Daniel Green, Arianna Ornaghi, Scott Nelson and Sebastian Fanelli.

Despite the importance of economists in helping frame and develop this dissertation, it certainly would not have been the same without the support and friendship of my friends from college (as well as some who technically might not have attended the University of Chicago but are by now as close as those who did). I was lucky to know many wonderful people in Boston, including Michael Newman, Jonathan Rodrigues, Bianca Blakesley, Nakul Singh, Hannah Fine, Stacey Kirkpatrick, Libby Bova, Grace Evans, Elissa Kaufman, Danielle Boucher, Surya Rajan, Andrew Kalweit, Blaire Byg, Ali Greatsinger, Katherine Rittenhouse, Spencer Tomarken, Carl Brozek, Joanna Spyra, Sam Moy and Sharon Tai. Though they did not live in Boston, I was lucky to be able to keep in touch with Michael Wong, Jay Garlapati, Mikhail Levin, Andy Dahl, Jenni Zelenty, Erik Strand, Brian Dressner, Rachel Meng, Nick Ramsey, Paige North, Maryclare Griffin, Persephone Ma, Dylan Piatt, Lou Smith, Alanna

Warner, Jason Weinberg, Helen Beilinson, David Crespo, Ben Thevathasan, Tali Datskovsky, Emily Ho, Augusta Gudeman, Michele Masiello, Emily Casey, Peter Ianakiev, Youssef Kalad, and Colin Cage and still consider them all wonderful friends.

Finally, I owe my work ethic and my love of learning to my parents most of all, who instilled in me the tools I needed from a young age. This dissertation is dedicated to them, for all the help they provided me and having unflinchingly supported me through all the stress that spilled over to them. The support of my brother Brian and sister-in-law Jackie were indispensible at multiple times during the course of graduate school. This process would have been immeasurably more difficult without them being close by. And I was lucky to have the love, support, and well wishes of my grandparents, cousins, aunts and uncles. I am very lucky to have such an amazing family.

To my parents

Introduction

The first chapter of this dissertation studies the question of optimal dynamic monopoly pricing under informational robustness. Consumers may be unsure of their willingness-to-pay for a product if they are unfamiliar with some of its features or have never made a similar purchase before. How does this possibility influence optimal pricing? To answer this question, we introduce a dynamic pricing model where buyers have the ability to learn about their values for a product over time. A seller commits to a pricing strategy, while buyers arrive exogenously and decide when to make a one-time purchase. The seller does not know how each buyer learns about his value for the product, and seeks to maximize profits against the worst-case information arrival processes. We show that a constant price path delivers the optimal profit, which is also the optimal profit in an environment where buyers cannot delay. We discuss the role of price-dependent information for this result, and consider an extension with common values and informational externalities.

The second chapter develops a model of costly information acquisition, focusing on an application to scientific research. When dimensions of an experiment that may bias a scientific result are not verified, scientists are incentivized to make their experiments more susceptible to false positives, even though they obtain higher surplus from more informative experiments. On the other hand, non-transparency can induce a scientist to undertake a costlier but more informative experiment if it also enables her to commit to acting scrupulously. Using this insight, this paper demonstrates the general existence of non-degenerate experiment costs such that greater transparency in scientific methodology results in research choices that are worse for those interested in the results.

The third chapter studies a model of delegated project choices with multiple agents, motivated by an application to the procurement of new technologies. These allocations are often done prior to the total resolution of uncertainty regarding ultimate viability. A common method of allocation involves prototyping—demonstrating some preliminary version of the project with the intention of bringing it to completion. This considers the value of competition in these settings, relating the results to challenges faced by policymakers. When agents share the same ordinal preferences over their projects as the principal (for instance, if surplus is ultimately divided according to Nash bargaining between the principal and selected agent), an optimal mechanism need not feature stochastic project choice when randomizations across agents are feasible. Such mechanisms may be suboptimal without the alignment of the principal's and agent's incentives. The resulting policy requires commitment whenever randomization is involved, despite the optimality of deterministic mechanisms in the single agent case.

Chapter 1

Informational Robustness in Intertemporal Pricing [1]

1.1 Introduction

Suppose a monopolist has invented a new durable product, and is deciding how to set prices over time to maximize profit. Consulting the literature on intertemporal pricing,[2] the monopolist would find that keeping the price fixed (at the single-period profit maximizing price) is an optimal strategy when consumers understand the product perfectly (provided willingness-to-pay does not vary over time). But a wrinkle arises if consumers may learn something that influences how much they like the product after pricing decisions have been made, a salient issue since the monopolist's product is completely new. For example, when the Apple Watch, Amazon Echo, and Google Glass were released, most consumers had little prior experience to inform their willingness-to-pay. In such a situation, the monopolist might suspect that the buyers' purchasing decisions will depend on the available information–e.g., journalist reviews about the product–which may in turn depend on the pricing decisions.

[1] Co-authored with Xiaosheng Mu

[2] E.g., Stokey (1979), Bulow (1981), Conlisk, Gerstner and Sobel (1984), among others. These papers show that a seller with commitment does not benefit from choosing lower prices in later periods.

The potential for information arrival presents a challenge to the monopolist's problem.

In isolation, components of this setting have been studied extensively. The literature on advertising, for instance, has considered the value of information for new products, treating it as given that there is some information that would inform consumers of their willingness-to-pay (see Bagwell (2007) for a thorough discussion of informative advertising). In the intertemporal pricing literature, Stokey (1979) recognized that willingness-to-pay may change over time, and that such changes can influence the optimal pricing strategy. And other papers on intertemporal pricing, such as Biehl (2001) and Deb (2014), have used exogenous learning by consumers to motivate their studies of stochastic changes in buyer values.

Despite this apparent interest, we are not aware of any papers that study dynamic pricing while modeling information arrival explicitly. We suspect one major reason for this absence relates to technical difficulties. Buyers' purchasing decisions depend on the value of information, something that is complicated in static environments and (as far as we are aware) intractable in most general dynamic environments. While Deb (2014) and Garrett (2016) restore tractability by considering specific evolution of buyer values, the stochastic processes they consider violate the martingale condition imposed by Bayesian updating. Their approaches are suitable for studying settings with taste shocks, but they do not fully capture learning. So the question of how to price optimally in the face of information arrival is left unanswered.

We introduce a model of intertemporal pricing that incorporates dynamic information arrival, and proceed to demonstrate a benchmark result on the optimality of constant price paths. To do this, we adopt the approach of the active literature on robust mechanism design. A seller commits to a pricing strategy, while buyers observe signals of their values, possibly over time, each according to some information structure (or more precisely, information arrival process). We assume that the seller does not know any part of the information arrival processes that inform the buyers of their values,[3] and is concerned with the worst possible

[3]While we assume each buyer knows her entire information arrival process, we show in Appendix A.4.2

information structures given the pricing decisions. One justification for this worst-case analysis is that the seller may want to guarantee a good outcome, no matter what the information structures actually are.[4] For our application, another justification would be that an adversary (e.g. a competitor or disgruntled journalists) may be interested in minimizing the seller's profit. If the firm did not have total control over what information consumers might have access to, our framework would be appropriate.[5] As for the commitment assumption, introducing it circumvents issues related to the Coase conjecture. Without seller commitment, this result implies that the worst-case is approximately achieved when buyers know their values, delivering profit to the seller equal to the minimum buyer valuation.[6]

Our first result is that a longer time horizon does not increase the amount of profit the seller can obtain from each buyer. One explanation is as follows: in each period, the adversary could release information that minimizes the profit in that period. Doing so would make the seller's problem separable across time, eliminating potential gains from decreasing prices. This intuition is incomplete, because the worst-case information structures for different periods need not be consistent, in the sense that past information may prevent the adversary from minimizing profits in the future. While this feature makes it difficult to find the exact worst case for arbitrary price paths, we use *partitional* information arrival processes to demonstrate how the adversary can hold the seller to a profit no larger than the single period benchmark. These information arrival processes involve the buyer learning whether his value is above or below a given threshold, with this threshold declining over time. As long as the seller desires robustness against this restrictive (yet intuitive) class of information arrival processes, our results would continue to hold.

---

that the results continue to hold if buyers face maxmin uncertainty over what information they will receive in the future.

[4]A more complete discussion of this justification can be found in the robust mechanism design literature, in particular: Chung and Ely (2007), Frankel (2014), Yamashita (2015), Bergemann, Brooks and Morris (2017), Carroll (2015, 2017).

[5]In Appendix A.4.2, we show that our solution also emerges in an alternative game where the information disclosure policy is set to help the buyer (rather than hurt the seller). This interpretation is inspired by (though distinct from) Roesler and Szentes (2017), which we discuss in depth.

[6]See Section 1.3.1 for a more complete discussion of the commitment assumption.

While the above argument shows that selling only once (at the single-period optimal price) is an optimal strategy with only a single buyer, this pricing strategy forgoes potential future profit when multiple buyers with i.i.d. values arrive over time. In the classic setting with known values, a constant price path maximizes the profit obtained from each arriving buyer, who either buys immediately upon arrival or not at all. This argument does not extend to our problem, since nature[7] can induce delay by promising to reveal information to the buyer in the future. Such delay could be costly for the seller, due to discounting. However, we show that as nature attempts to convince the buyer to delay her purchase, it must also promise a greater probability of purchase to satisfy the buyer's incentives. It turns out that, from the seller's perspective, the cost of delayed sale is always offset by the increased probability of sale. We thus show that a constant price path ensures the greatest worst-case profit, equal to the profit when buyers can only possibly buy upon arrival. We refer to this profit guarantee as the *no-delay guarantee*.

Together, this analysis delivers a result qualitatively similar to one that has been shown under known values (e.g., Stokey (1979)): The seller's optimal strategy is to hold the price fixed at the single-period optimal price, and (in the worst-case) buyers purchase either immediately or never. This holds even though the single-period optimum in our problem is different from Stokey (1979) due to buyer learning.

A crucial assumption in our main model is that buyer information in each period can depend on the entire history of realized prices. In Section 1.7, we consider several variants of the model, which allow for less interaction between prices and information. These alternative setups generalize the single-period models studied by Du (2018) and Roesler and Szentes (2017). Though their (single-period) profit guarantee is typically higher than ours, we discuss conditions on the information arrival process that ensure the no-delay guarantee is still achievable.

Section 1.8 shows that the difference between our profit guarantee and that of Roesler and Szentes (2017) and Du (2018) disappears when players are patient and informational

[7]It is convenient to think of "nature" choosing the information arrival process to hurt the seller.

externalities are present. Specifically, we consider a variant of the model where the value for the product is common across buyers and information is publicly observed. In this alternative model, whether or not information can depend on the current period price does not influence the optimal profit guarantee when all players are sufficiently patient.

We begin by reviewing the literature, and then proceed to present the main model. The one period benchmark is studied in Section 1.4, and we show that intertemporal incentives do not help the seller in Section 1.5. Using this result, we demonstrate that constant price paths are optimal in Section 1.6. The remaining sections discuss our timing assumptions, informational externalities and other modifications. Section 1.10 concludes.

1.2 Literature Review

This paper is part of a large literature that studies pricing under robustness concerns, where the designer may be unsure of some parameter of the buyer's problem. Informational robustness is a special case, and one that has been studied in static settings. The most similar to our one-period model are Roesler and Szentes (2017) and Du (2018). Both papers consider a setting like ours, where the buyer's value comes from some commonly known distribution, but where the seller does not know the information structure that informs the buyer of her value.[8] Taken together, these papers characterize the seller's maxmin pricing policy and nature's minmax information structure in the *static* zero-sum game between them.[9] The one-period version of our model differs from these papers, since we assume that nature can reveal information depending on the *realized* price the buyer faces (see Section 1.3.1 for further discussion). Moreover, our paper is primarily concerned with dynamics, which is absent from Roesler and Szentes (2017) and Du (2018).

[8]Du (2018) extends the analysis to a one-period, many-buyer common value auction environment. He constructs a class of mechanisms that extracts full surplus when the number of buyers grows to infinity, despite the presence of informational uncertainty. This is solved in the special case of two buyers and two value types by Bergemann, Brooks and Morris (2016), and in the general case by Brooks and Du (2018).

[9]Roesler and Szentes (2017) actually motivate their model as one where the buyer chooses the information structure; they show that this solution also minimizes the seller's profit. See Appendix A.4.2 for a related interpretation of our model.

Other papers have considered the case where the value distribution itself is unknown to the seller. For instance, Carrasco et al. (2017) consider a seller who does not know the distribution of the buyer's value, but who may know some of its moments. If the distribution has two-point support, our one-period model becomes a special case of Carrasco et al. (2017) in which the seller knows the support as well as the expected value. But in general, even in the static setting, assuming a prior distribution constrains the possible posterior distributions nature can induce beyond any set of moment conditions. This point is elaborated on in Section 1.9.2.

In our model, nature being able to condition on realized prices is sufficient to eliminate any gains to randomization (even if the randomization is to be done in the future). This may be reminiscent of Bergemann and Schlag (2011), who show (in a one-period model) that a deterministic price is optimal when the seller only knows the true value distribution to be in some neighborhood of distributions.[10] However, the reasoning in Bergemann and Schlag (2011) is that a single choice by nature yields worst-case profit for all prices. This is not true in our setting, but we are able to construct an information structure for every pricing strategy that shows randomization does not have benefits.

While most of this literature is static, some papers have studied dynamic pricing where the seller does not know the value distribution. Handel and Misra (2014) allow for multiple purchases, while Caldentey, Liu, Lobel (2016), Liu (2016) and Chen and Farias (2016) consider the case of durable goods. As discussed above, information arrival restricts *how* the value evolves, and rules out the cases considered in the literature. In addition, these papers look at different seller objectives; the first three study regret minimization, whereas the last one looks at a particular mechanism that approximates the optimum.

The literature on robust mechanism design has become popular in recent years in part due to its ability to provide foundations for the optimality of simple mechanisms, which tend to be observed in practice. For instance, Carroll (2017) shows how uncertainty over

[10]Their result applies to maxmin profit as in our model. The authors also show that if the seller's objective is instead to minimize regret, then random prices do better.

the correlation between a buyer's demand for different goods leads to the seller pricing the goods independently.[11] In Carroll (2015), uncertainty over the mapping from an agent's actions into output leads to the principal aligning the agent's compensation directly with output. In Frankel (2014), similar alignment arises when there is uncertainty over the agent's bias, and Yamashita (2015) shows how uncertainty about bidders' higher order beliefs may favor second price auctions even with interdependent values. At the moment, however, this literature has had much less to say about dynamic environments. Important exceptions are Penta (2015), who considers the implementation of social choice functions in dynamic settings, and Chassang (2013), who shows how dynamics allow a principal to approximate robust contracts which may be infeasible in the presence of liability constraints. These are both rather different from our setting, and we suspect there is much work left to be done in this area.

Several intertemporal pricing papers allow for the value to change over time without explicitly modeling information arrival (absent robustness concerns). Stokey (1979) assumes the value changes deterministically given the initial type. Deb (2014) assumes the value is independently redrawn upon Poisson shocks. For Garrett (2016), the value follows a two-type Markov-switching process. As mentioned above, these papers do not impose the martingale condition for expectations. We are not aware of how to determine buyer purchasing behavior under an arbitrary information arrival process. But the maxmin objective allows us to focus on simple and intuitive information structures, making the buyer's problem tractable.

Finally, it is well known that the literature on informational robustness is related to the literature on information design, which has also recently begun to study dynamics (see Ely (2017) and Ely, Frankel and Kamenica (2015)). While we are ultimately concerned with pricing strategies, this connection is relevant because we describe how a receiver's (buyer)

[11] The general link between dynamic allocations and multi-dimensional screening has been long noted in Bayesian settings (see, for instance, Pavan Segal and Toikka (2014) for a discussion of this point). While it is interesting that we obtain a result that is similar to his, we note that our focus on information arrival and single-object purchase are significant differences.

behavior varies depending on how a sender (nature) chooses the information structure. Several of our results (in particular, the proof of Lemma 2) bear resemblance to this literature, and they may be of interest outside of our setting.

1.3 Model

A seller sells a durable good at times $t = 1, 2, \ldots, T$, where $T \leq \infty$. At each time t, a single buyer arrives and decides if and when to buy the object.[12] All parties discount the future at rate δ. The product is costless for the seller to produce,[13] while each buyer has unit demand. Throughout what follows, we let t denote calendar time, and let a index a buyer's arrival time. Each buyer has an independently drawn discounted lifetime utility from purchasing the object. We let v denote some unspecified buyer's value, and assume that each buyer's value is drawn from a distribution F supported on $\mathbb{R}_+$, with $0 < \mathbb{E}[v] < \infty$. We let $\underline{v}$ denote the minimum value in the support of F. The distribution F is fixed and common knowledge, and buyer values for the object do not change over time.

However, the buyers do not directly know their v; instead, they learn about it through signals they obtain over time, via some information structure. To be precise, a *dynamic information structure* $\mathcal{I}_a$ for a buyer arriving at time a is:

- A set of possible signals for every time t after a, i.e., a sequence $(S_t)_{t=a}^{T}$, and
- Probability distributions given by $I_{a,t} : R_+ \times S_a^{t-1} \times P^t \to \Delta(S_t)$, for all t with $a \leq t \leq T$.

Without loss of generality, we assume that all buyers are endowed with the same signal sets S_t, although each one privately observes any particular signal realization. Note that the buyer observes signal realization s_t at time t, whose distribution depends on (their own) true value

[12]Our basic analysis is unchanged if the number of arriving buyers varies over time, or is stochastic, provided the value distribution is fixed.

[13]Introducing a cost of c per unit does not change our results: it is as if the value distribution F were "shifted down" by c, and the buyer might have a negative value. The transformed distribution G in Definition 1 below would also be shifted down by c.

$v \in R_+$, the history of (their own) previous signal realizations $s_a^{t-1} = (s_a, s_{a+1}, \ldots, s_{t-1}) \in S_a^{t-1}$, as well as the history of *all previous and current* prices $p^t = (p_1, p_2, \ldots, p_t) \in P^t$. In particular, this definition allows for information structures to display history dependence.[14]

The timing of the model is as follows. At time 0, the seller commits to a pricing strategy σ, which is a distribution over possible price paths $p^T = (p_t)_{t=1}^T$. We allow $p_t = \infty$ to mean that the seller refuses to sell in period t. Note that the price the seller posts at time t must be the same for all buyers that have the ability to buy in that period. After the seller chooses the strategy, nature chooses a dynamic information structure for each buyer. In each period $t \geq 1$, the price in that period p_t is realized according to $\sigma(p_t \mid p^{t-1})$. A buyer arriving at time a with true value v observes the signal s_t with probability $I_{a,t}(s_t \mid v, s_a^{t-1}, p^t)$ and decides whether or not to purchase the product (and if so, when).

Given the pricing strategy σ and the information structure $\mathcal{I}_a$, the buyer arriving at time a faces an optimal stopping problem. Specifically, they choose a stopping time τ_a^* adapted to the joint process of prices and signals, so as to maximize the expected discounted value less price:

$$\tau_a^* \in \underset{\tau}{\operatorname{argmax}}\ \mathbb{E}\left[\delta^{\tau-a}(\mathbb{E}[v|s_a^\tau, p^\tau] - p_\tau)\right].$$

The inner expectation $\mathbb{E}[v|s_a^\tau, p^\tau]$ represents the buyer's expected value conditional on realized prices and signals up to and including period τ. The outer expectation is taken with respect to the evolution of prices and signals. We note that the stopping time τ_a is allowed to take any positive integer value $\leq T$, or $\tau_a = \infty$ to mean the buyer never buys.

The seller evaluates payoffs as if the information structure chosen by nature were the worst possible, given his pricing strategy σ and buyer's optimizing behavior. Hence the seller's payoff is:

$$\sup_{\sigma \in \Delta(p^T)} \inf_{(\mathcal{I}_a), (\tau_a^*)} \sum_{a=1}^{T} \mathbb{E}[\delta^{\tau_a^* - a} p_{\tau_a^*}] \text{ s.t. } \tau_a^* \text{ is optimal given } \sigma \text{ and } \mathcal{I}_a, \forall a.$$

Note that when a buyer faces indifference, ties are broken against the seller. Breaking

[14] To avoid measurability issues, we assume each signal set S_t is at most countably infinite. All information structures in our analysis have this property.

indifference in favor of the seller would not change our results, but would add cumbersome details.[15]

1.3.1 Discussion of Assumptions

Several of our assumptions are worth commenting on. First, following the robust mechanism design literature, we assume that the buyer has perfect knowledge of the information structure whereas the seller does not. More precisely, each buyer knows the information structure, and is Bayesian about what information will be received in the future. In contrast, the seller is uncertain about the information structure itself. Our interpretation is that the buyers understands what information they will have access to; for instance, someone may always rely upon some product review website and hence know very well how to interpret the reviews. The seller, on the other hand, knows that there are many possible ways buyers can learn, and wants to do well against all these possibilities. In Section 1.9.1, we will show that our results extend even if the seller knows the buyer begins with extra prior information (say, through advertising). Thus, a deterministic constant price path remains optimal when nature is constrained to provide some particular information (but could provide more) in the first period. In Appendix A.4.2, we also show that, as long as the buyer is uncertainty averse and knows how to interpret all signals they have received, the worst case for the seller involves a Bayesian buyer. Our results only require that the buyers know what information they receive at any given time.

Second, we assume that the value distribution is common knowledge. This restriction is for simplicity, allowing us to focus on information arrival and learning. The assumption also enables us to compare our results to the classic literature on intertemporal pricing. In fact, the classic setting where the buyer knows her value can be seen as an extreme case of our extended model in Section 1.9.1.

Third, we assume that the information structure for a buyer arriving at a only depends

[15] When ties are broken against the seller, it follows from our analysis that the sup inf is achieved as max min. This would not be true if ties were broken in favor of the seller.

on his value, his signal history and the price history. In principle, one may want the information structure to depend on more variables, such as the purchasing history or the signals and values of other buyers. However, because of our worst case objective and the i.i.d. assumption, allowing for nature to condition on more variables would not hurt the seller further. In contrast, these variables are important in Section 1.8, when we allow common values and public information.

Fourth, we assume that the seller commits to a pricing strategy that the buyer observes. This avoids certain technical difficulties related to formalizing (seller) learning under ambiguity (see Epstein and Schneider (2007)). But in practice, firms like Amazon and Apple are widely followed by consumers and industry experts, meaning that they are typically able to credibly announce (and stick to) consistent pricing strategies. And while some strategies may be difficult for a seller to commit to, constant price paths are significantly simpler to implement since deviations are straightforward to detect. On the other hand, we *restrict* the seller to using pricing mechanisms, and rule out, for instance, mechanisms that randomly allocate the object as a function of reports. We view this as a restriction on the environment, and one that tends to be quite common in our applications.

Finally, our key timing assumption is that the information structure in each period is determined *after* the price for that period has been realized. As in the literature review, if the information structure is determined *before* the price is realized, then the one-period optimal seller strategy would follow from Roesler and Szentes (2017) and Du (2018). The question of timing is more delicate under dynamics; should a buyer's second period information depend on the first period price they observed? What about buyers that arrive later? These questions motivate the analysis in Section 1.7, which discusses these issues in more depth. In any event, we believe that information could depend (at least somewhat) on price in practice. When shopping online, a consumer's information about a product may depend on how prominently it is displayed in the search results. If the buyer sorts products by how expensive they are, then the information structure will depend on the realized price.

1.4 Single Period Analysis

We start with the case where the seller does not worry about intertemporal incentives. For simplicity, we do this by taking $T = 1$, although the results are identical if buyers are myopic or could only purchase upon arrival. To solve this problem, we define a transformed distribution of F. For expositional simplicity, the following definition assumes F is continuous. All of our results in this paper extend to the discrete case, though the general definition requires additional care and is relegated to Appendix A.1.

Definition 1. *Given a continuous distribution F, the transformed distribution $G = P(F)$ is defined as follows. For $y \in \mathbb{R}_+$, let $L(y)$ denote the conditional expectation of $v \sim F$ given $v \leq y$. Then G is the distribution of $L(y)$ when y is drawn according to F. We call G the "pressed" version of F, and refer to the mapping P as "pressing."*

The pressed distribution G is useful because for any (realized) price p, nature can only ensure that the object remains unsold with probability $G(p)$. This holds since the worst-case information structure has the property that a buyer who does not buy has expected value exactly p. To see why, consider an information structure where the buyer's belief following a recommendation to not buy is $v_N < p$ and the recommendation following a recommendation to buy is $v_B(> p)$. Then nature could, with some small probability $\varepsilon > 0$, give the recommendation to not buy whenever buy would have been recommended, hurting the seller. The buyer would have a higher belief following a recommendation to not buy, but would still follow it if ε were sufficiently small. This logic holds as long as $v_N < p$.

In fact, the worst-case information structure following a price p is a *partition* with a threshold that induces a belief p following the recommendation not to buy. One can show (e.g., Kolotilin (2015)) that partitional information structures minimize the probability the buyer is recommended to buy, whenever the belief following the recommendation to not buy must be some fixed value (in our case, p). This observation allows us to show that the worst-case information structure involves telling the buyer whether her value is above or below $F^{-1}(G(p))$, making $1 - G(p)$ the probability of sale.

These remarks give us the following proposition:

Proposition 1. *In the one-period model, a maxmin optimal pricing strategy is to charge a deterministic price p^* that solves the following maximization problem:*

$$p^* \in \underset{p}{\operatorname{argmax}}\ p(1 - G(p)). \tag{1.1}$$

For future reference, we call p^* the one-period maxmin price and similarly $\Pi^* = \max_p p(1 - G(p))$ the one-period maxmin profit.

It is worth comparing the optimization problem (1.1) to the standard model without informational uncertainty. If the buyer knew her value, the seller would maximize $p(1 - F(p))$. In our setting, the difference is that the transformed distribution G takes the place of F, which will be useful for the analysis in later sections. The following example illustrates:

Example 1. *Let $v \sim$ Uniform[0,1], so that $G(p) = 2p$. Then $p^* = \frac{1}{4}$ and $\Pi^* = \frac{1}{8}$. With only one period to sell the object, the seller charges a deterministic price $1/4$. In response, nature chooses an information structure that tells the buyer whether or not $v > 1/2$.*

In Example 1, relative to the case where the buyer knows her value, the seller charges a lower price and obtains a lower profit under informational uncertainty. In Appendix A.4.2, we show that this comparative static holds generally.

Finally, also note that there are other information structures which induce the same worst-case profit for the seller. For example, the buyer could be told her value exactly if it is above the threshold, since she will still buy. However, any worst case information structure involves the buyer being told if her value is *below* the threshold (i.e., the lowest element of the partition cannot be refined further on a set of positive measure).

1.5 Intertemporal Incentives Do Not Help

In this section we present our first main result, that having multiple periods to sell does not allow the seller to extract more surplus from each buyer.[16] Stokey (1979) demonstrated that this result holds when buyers know their values, provided they do not change over time. On the other hand, she also demonstrated that if values do change over time, letting buyers delay purchase could enable a seller to obtain higher profits by facilitating price discrimination.[17] One may wonder whether information arrival, which affects the buyers' value over time, could similarly make price discrimination worthwhile.

However, it turns out for worst case information structures, these concerns do not arise. For simplicity, we focus on the case where there is a single buyer at time 1, since the argument readily extends to the case where buyers arrive at every time. With only the first buyer, the seller could always sell exclusively in the first period, the one-period profit Π^* forms a lower bound on the seller's maxmin profit from this buyer. To show that Π^* is also an upper bound, we explicitly construct a dynamic information structure for any pricing strategy, such that the seller's profit under this information structure decomposes into a convex combination of one-period profits. Our proof takes advantage of the partitional form of worst-case information structures from the single period problem:

Proposition 2. *For any pricing strategy $\sigma \in \Delta(p^T)$, there is a dynamic information structure $\mathcal{I}$ and a corresponding optimal stopping time τ^* that lead to expected (undiscounted, per-buyer) profit no more than Π^*. So, for a single buyer, the seller's maxmin profit against all dynamic information structures is Π^*, irrespective of the time horizon T and the discount factor δ.*

We will present the proof of this proposition under the assumption that the seller charges

[16] We highlight that the dynamics of information arrival are crucial for this result. For instance, suppose the seller knew that information would *not* be released in some period t. Then he could sell exclusively in that period and (by charging random prices) obtain the Roesler and Szentes (2017) profit level, which is generally higher than Π^* (see Section 1.7 for details). For δ sufficiently close to 1, this pricing strategy does better than a constant price path.

[17] It is interesting to note in our worst case information structures, buyers who do not buy actually *do* have a positive continuation value, even though this need not hold for arbitrary (non-worst case) partitional information arrival processes.

a deterministic price path $(p_t)_{t=1}^T$. This is *not* without loss, because random prices in the future may make it more difficult for nature to choose an information structure in the current period that minimizes profit. However, our argument does extend to random prices and shows that randomization does not help the seller. We discuss this after the (more transparent) proof for deterministic prices.

Let us first review the sorting argument when the buyer knows her value. In this case, given a price path $(p_t)_{t=1}^T$, we can find time periods $1 \leq t_1 < t_2 < \cdots \leq T$ and value cutoffs $w_{t_1} > w_{t_2} > \cdots \geq 0$, such that the buyer with $v \in [w_{t_j}, w_{t_{j-1}}]$ optimally buys in period t_j (see e.g. Stokey (1979)). This implies that in period t_j, the object is sold with probability $F(w_{t_{j-1}}) - F(w_{t_j})$.

Inspired by the one-period problem, we construct an information structure under which in period t_j, the object is sold with probability $G(w_{t_{j-1}}) - G(w_{t_j})$ (that is, where G replaces F). The following information structure $\mathcal{I}$ has this property:

- In each period t_j, the buyer is told whether or not her value is in the lowest $G(w_{t_j})$-percentile.
- In all other periods, no information is revealed.

This information structure is similar to the one period problem, in that a buyer is told whether her value is above a given threshold. In the dynamic setting, this threshold is now declining over time. We refer to these information structures as *partitional* information arrival processes, since different signal realizations partition the support of the buyer's value distribution into disjoint intervals. Note that the thresholds are chosen to make the buyer indifferent between purchasing *and continuing without further information*. The buyer therefore prefers to delay purchase when her value is below the threshold. On the other hand, a buyer whose value is above the threshold does not expect to receive further information, and hence purchases immediately. These observations are summarized in the following lemma:

Lemma 1. *Given prices $(p_t)_{t=1}^T$ and the information structure $\mathcal{I}$ constructed above, an optimal stopping time τ^* involves the buyer buying in the first period t_j when she is told her value is **not** in the lowest $G(w_{t_j})$-percentile.*

The proof of this lemma can be found in Appendix A.1, where we actually prove a more general result for random prices.

Using this lemma, we can now prove Proposition 2 by computing the seller's profit under the information structure $\mathcal{I}$ and the stopping time τ^*:

Proof of Proposition 2 for Deterministic Prices. Since the buyer with true value v in the percentile range $(G(w_{t_j}), G(w_{t_{j-1}})]$ buys in period t_j, the seller's discounted profit is given by (assuming $T = \infty$):

$$\begin{aligned} \Pi &= \sum_{j \geq 1} \delta^{t_j - 1} p_{t_j} \cdot \Big(G(w_{t_{j-1}}) - G(w_{t_j}) \Big) \\ &= \sum_{j \geq 1} (\delta^{t_j - 1} p_{t_j} - \delta^{t_{j+1} - 1} p_{t_{j+1}}) \cdot (1 - G(w_{t_j})) \\ &= \sum_{j \geq 1} (\delta^{t_j - 1} - \delta^{t_{j+1} - 1}) w_{t_j} \cdot (1 - G(w_{t_j})) \\ &\leq \delta^{t_1 - 1} \cdot \Pi^*, \end{aligned} \tag{1.2}$$

where the second line is by Abel summation,[18] the third line is by w_{t_j}'s indifference between buying in period t_j or t_{j+1}, and the last inequality uses $w_{t_j}(1 - G(w_{t_j})) \leq \Pi^*, \forall j$. For finite horizon T, the proof proceeds along the same lines except for a minor modification to Abel summation. ■

Relative to the potential complexity of arbitrary information arrival processes, we find it noteworthy that the information structures constructed here are reasonably intuitive: Consumers buy when they find out that their value is above some (price contingent) threshold. Intertemporal pricing cannot help the seller as long as she is concerned (at least) with these particular information arrival processes.

[18] Abel summation says that $\sum_{j \geq 1} a_j b_j = \sum_{j \geq 1} \Big((a_j - a_{j+1}) \sum_{i=1}^j b_i \Big)$ for any two sequences $\{a_j\}_{j=1}^\infty, \{b_j\}_{j=1}^\infty$ such that $a_j \to 0$ and $\sum_{i=1}^j b_i$ is bounded. We take $a_j = \delta^{t_j - 1} p_{t_j}$ and $b_j = G(w_{t_{j-1}}) - G(w_{t_j})$.

Despite the appeal of the analogy to the known value case, it is worth noting that for an arbitrary declining price path, these information structures we construct in the proof above *may not* be the worst-case. The following example illustrates (we use a discrete value distribution for simplicity):

Example 2. *Let $T = 2$, $v \in \{0, 1\}$ with $\mathbb{P}[v = 1] = 1/2$ and $\delta = 1/2$. Suppose the seller were to use a price path $p_1 = 11/40$ and $p_2 = 1/10$. Since a buyer would be indifferent between purchase and delay with a true value of $\frac{9}{20}$, the information structure constructed in Lemma 1 applied to this example induces posterior expected value $\frac{9}{20}$ when the buyer is recommended to not purchase in the first period, and expected value p_2 when recommended to not purchase in the second period. One can show that the (overall) expected profit is:*[19]

$$p_1 \cdot \frac{1}{11} + (\delta p_2) \cdot \left(1 - \frac{1}{11}\right) \left(\frac{7}{18}\right) \approx 0.0427 < 0.0858 \approx \Pi^*.$$

Now suppose that instead, nature were to provide no information in the first period and reveal the value perfectly in the second period. Note that the buyer would be willing to delay, since

$$\mathbb{E}[v] - p_1 \leq \delta \cdot \mathbb{P}[v = 1] \, (1 - p_2) \, .$$

In fact, equality holds. Under this information structure, the seller's profit is therefore $\delta p_2 \mathbb{P}[v = 1] = \frac{1}{40} = 0.025 < 0.0427$.

The important feature of the example is that in the second period, the buyer *strictly* prefers following the recommendation they are given to disobeying it. Such an information structure creates option value, and potentially hurts the seller by inducing delay in dynamic settings. While we can show that the above information structure is indeed worst case in this example, finding the worst-case information structures against a given price path seems challenging in general, and is not necessary for our main result on the optimality of constant prices.

We note that random prices introduce a technical difficulty in applying the sorting

[19] If the probability of being recommended to buy in period t is r_t, we have $\frac{1}{2} = r_1 + \frac{9}{20}(1 - r_1)$ and $\frac{9}{20} = r_2 + \frac{1}{10}(1 - r_2)$. These equations give $r_1 = \frac{1}{11}$ and $r_2 = \frac{7}{18}$.

argument directly. Since the threshold values w_{t_j} depend on both the realized price and the distribution of future prices, they are in general random variables. More problematically, these thresholds may be non-monotonic if they are defined using the buyer's indifference condition. If such non-monotonicity occurs, we will not be able to express the seller's discounted profit as a convex sum of one-period profits, and the profit bound will not be valid.

In Appendix A.1 we show that the intuition from the deterministic case still works when prices can be random, but we develop additional technical tools in order to generalize the construction appropriately. Specifically, we modify the relevant indifference thresholds so that they are *forced* to be decreasing. To be precise, we define v_t to be the smallest value (in the known-value case) that is indifferent between buying in period t at price p_t and optimally stopping in the future, and *then* let $w_t = \min\{v_1, v_2, \ldots, v_t\}$. We think of this as keeping track of the "binding" thresholds, above which all consumers have already bought. This circumvents the monotonicity issue that arises with randomizations, so that we can use the redefined w_ts in our specification of the (otherwise identical) dynamic information structure. The rest of the proof is as before, with the assistance of Lemma 9, which expresses the price at any period as the expectation of a convex sum of the present and future w_ts.[20] Proposition 2 thus continues to hold for random prices.

1.6 Optimality of Constant Prices

We now demonstrate the optimality of constant price paths. By Proposition 2, the seller's discounted profit from the buyer arriving at time a is bounded above by $\delta^{a-1} \cdot \Pi^*$. This gives us an upper bound for the seller's worst case profit. Furthermore, if the seller were able to set personalized prices (i.e., conditioning on the arrival time), this upper bound could be achieved by selling only once to each arriving buyer. We will show that the seller can achieve the same profit level by always charging p^*, without conditioning prices on the

[20] This identity replaces the indifference condition utilized to derive the third line of Equation (1.2).

arrival time.

Under known values, any arriving buyer facing a constant price path would buy immediately (if she were to buy at all), due to impatience. However, the promise of future information may induce the buyer to delay, even with constant prices. Nevertheless, in the following lemma, we show that against non-decreasing price paths, it is enough to consider information structures which only release information upon arrival:

Lemma 2. *In the multi-period model with one buyer, the seller can guarantee* Π^* *with any deterministic price path* $(p_t)_{t=1}^T$ *satisfying* $p^* = p_1 \leq p_t, \forall t$.

We present the intuition here and leave the formal proof to Appendix A.1. Fixing a non-decreasing price path and an arbitrary dynamic information structure nature could choose, we consider an alternative information structure that only gives a recommendation to the buyer (to purchase or not) when they arrive. The probability that the buyer is recommended to purchase at time 1 in this replacement information structure leaves the *discounted* probability of sale unchanged. In other words, we "push and discount" nature's recommendation to the buyer's arrival time.

The proof shows that for non-decreasing prices, the buyer would follow the recommendation of this replacement, while the seller's profit is weakly decreased. Since the seller receives at least Π^* under any information structure that releases information only in the first period, we obtain the lemma. Note that Example 2 demonstrates this argument relies upon non-decreasing prices.

Armed with this lemma, we can show our main result of the paper. The proof is straightforward given our results:

Theorem 1. *The seller can guarantee* $\Pi^* \cdot \frac{1-\delta^T}{1-\delta}$ *with a constant price path charging* p^* *in every period. This deterministic pricing strategy is optimal, and it is uniquely optimal whenever the one-period maxmin price* p^* *is unique.*

Facing a constant price path, a worst-case dynamic information structure simply involves giving each buyer the same information structure they would have obtained with only a single period. This completes our analysis of the baseline model.

1.7 Timing

This section analyzes the implications of our assumption regarding the timing of information acquisition relative to pricing. This assumption is captured in how we define dynamic information structures, since we allow them to be contingent on all past prices as well as the current price, but not future prices. When $T = 1$, the benchmark where information cannot depend on price is studied in Roesler and Szentes (2017) and Du (2018), which together solve the seller's (and nature's) problem in this benchmark.[21] For completeness, we recall their result. To make the connection with our paper most clear, we impose as in these papers that the buyer's value distribution F is supported on $[0, 1]$. Roesler and Szentes (2017) observe that in choosing an information structure, nature is equivalently choosing a distribution $\tilde{F}$ of posterior expected values, such that F is a mean-preserving spread of $\tilde{F}$.[22] They solve for the worst-case distribution $\tilde{F}$ as summarized below:

Theorem 1 in Roesler and Szentes (2017). *For $0 \leq W \leq B \leq 1$, consider the following distribution that exhibits unit elasticity of demand (with a mass point at $x = B$):*

$$F_W^B(x) = \begin{cases} 0 & x \in [0, W) \\ 1 - \frac{W}{x} & x \in [W, B) \\ 1 & x \in [B, 1] \end{cases} \tag{1.3}$$

In the one-period zero-sum game between the seller and nature, an optimal strategy by nature is to induce posterior expected values given by the distribution F_W^B for some W, B, such that W is smallest possible subject to F being a mean-preserving spread of F_W^B.

---

[21] In principle, these may not be the only two benchmarks of interest, since one could also study cases where information interacts somewhat, but not arbitrarily, with the price. We do not do so here since we are not aware of any reasonable, compelling alternative restrictions.

[22] This equivalence is separately observed by Gentzkow and Kamenica (2016) in the context of Bayesian persuasion. These authors attribute the result to Rothschild and Stiglitz (1970).

The seller's optimal single-period profit guarantee is found by computing the smallest W defined above, which we denote by Π_{RSD}. Conversely, Du (2018) constructs a particular mechanism the seller can use to guarantee profit Π_{RSD} against any information structure. In Appendix A.4.1, we represent Du's "exponential mechanism" as an equivalent random price mechanism. We also note that $\Pi_{RSD} \geq \Pi^*$ in general, and in Appendix A.4.1 we characterize when the inequality is strict.

As alluded to in the Introduction and Section 1.3.1, specifying the role of prices in dynamic information structures is more subtle than for static information structures. Over time, there are many more ways for information to interact with price (or not). Our benchmark corresponds to the most cautious case, but other cases are worth commenting on as well (particularly for readers interested in how Roesler and Szentes (2017) and Du (2018) extend to dynamic settings).

In this Section, we re-define a *dynamic information structure* to be a sequence of signal sets $(S_t)_{t=1}^T$ and probability distributions $I_{a,t} : R_+ \times S_a^{t-1} \times P^{t-1} \to \Delta(S_t)$. The crucial distinction from our main model is that the signal s_t depends on previous prices p^{t-1} but not on the current price p_t. The seller chooses a pricing strategy that achieves maxmin profit against these information structures (and optimal stopping times of the buyer).

Before moving to the full model with arriving buyers, as a warm up we note that a more direct argument can be used in this setting to show a seller does no better with a longer horizon than our main model required (i.e., avoiding the constructions in Proposition 2).

Proposition 3. *Suppose there is a single buyer. For any time horizon T and any discount factor δ, the seller's maxmin profit when nature cannot condition on the current period price is given by Π_{RSD}.*

The reasoning is as follows: With multiple periods and a single buyer, the seller can guarantee Π_{RSD} by selling only once in the first period (using Du's mechanism). On the other hand, suppose nature provides the Roesler-Szentes information structure in the first period and no additional information in later periods. Then the seller faces a fixed distribution of values given by F_W^B. By Stokey (1979), selling only once is optimal against this distribution, and the seller's optimal profit is at most $W = \Pi_{RSD}$. This proves the result.

While both Proposition 2 and Proposition 3 show a longer selling horizon does not help the seller, here the argument is more direct due to the duality between Roesler-Szentes and Du. In other words, without price dependence and with only one buyer, the transformation of the distribution can be done without reference to future prices. But in our baseline seting, this does not hold, as evidence by the fact that the dynamic information structures in Proposition 2 must depend on the seller's future pricing strategy. With arriving buyers, however, withholding future sales is costly, as in our baseline model, and we again need to worry about the possibility the future sales provides a channel for information arrival to harm the seller. We consider three different cases; in the first two, the no-delay profit is still achievable, but in the third it is not.

1.7.1 Case One: Information *only* upon arrival, but possibly contingent on past (though not current) prices

Our first benchmark considers the subset of dynamic information structures defined prior to Proposition 3 which only provide each buyer with a single signal. Hence each buyer is endowed with a single probability distribution $I_a : R_+ \times P^{t-1} \to \Delta(S_a)$. These are dynamic in the sense that they respond to prices, but not in the sense that information arrives over time. Note that if the seller used a constant price path, nature can provide $F^B_{\Pi_{RSD}}$ and the worst-case partition to later buyers. The seller's profit from all buyers after the first is then no larger than Π^*, and since $\Pi^* < \Pi_{RSD}$ with $v \sim U[0,1]$, a constant price path would not achieve the no-delay guarantee.

However, by changing the way the seller randomizes over the price, the no-delay guarantee is achievable:

Theorem 2. *Suppose there are arriving buyers, and suppose each buyer only receives information once upon arrival (before the price realizes in that period). For any time horizon T and any discount factor δ, the seller has a pricing strategy that ensures profit at least Π_{RSD} from each buyer. Thus the* *seller's maxmin profit is* $\Pi_{RSD} \cdot \left(\frac{1-\delta^T}{1-\delta}\right)$.

The proof is based on a key lemma (Lemma 10 in Appendix A.2) relating the outcome under a static price distribution to that under a dynamic price distribution. This outcome-equivalence property enables us to construct a dynamic pricing strategy that replicates Du's mechanism for each arriving buyer, achieving Π_{RSD} as profit guarantee. Essentially, we consider randomization over *threshold indifference conditions* instead of randomization over price itself, and then use this to construct the dynamic random pricing strategy of the seller. This intuition—namely, to consider the buyer's indifference conditions instead of the prices—bears some resemblance to the intuition behind Proposition 2 and Lemma 9, although we are not aware of any formal equivalence between the two.

1.7.2 Case Two: Information cannot depend on prices at all

Next, we consider the case where we define a dynamic information structure to be *completely* price independent—that is, a probability distribution $I_{a,t} : R_+ \times S_a^{t-1} \to \Delta(S_t)$—but allow for buyers to obtain information over time. While this case is in some sense the polar opposite of what we study, the resulting optimum is remarkably similar. Specifically, an optimal strategy for the seller is to utilize a constant price path, but drawing the price path randomly according to a Du distribution instead of setting it equal to p^* (as opposed to in our main model). It turns out that again, intertemporal incentives disappear when this strategy is employed:

Theorem 3. *Suppose there are arriving buyers, and suppose that all information is independent of all realized prices (though may depend on the pricing strategy). For any time horizon T and any discount factor δ, the seller has a pricing strategy that ensures profit at least Π_{RSD} from each buyer. Thus the seller's maxmin profit is* $\Pi_{RSD} \cdot \left(\frac{1-\delta^T}{1-\delta}\right)$.

This theorem involves new techniques that may be of independent interest. Recall that, in allowing for random pricing strategies in Section 1.5, we defined cutoff values (forced to be decreasing) in two steps—first using the buyer's indifference condition, and then keeping track of the lowest realized value arising from this indifference condition. Intuitively, these were the relevant "binding thresholds," above which consumers would have already bought. Inspired by this technique, the proof of Theorem 3 introduces the dual definition of cutoff thresholds, namely *cutoff prices*. This approach is natural since the Du mechanism gives a natural guess for the worst case pricing strategy (circumventing the need to construct an information structure against an arbitrary price path as in Section 1.5). Having this guess is not enough, however, since we still need to consider a much larger set of information structures than in Du (2018), and in particular cannot directly appeal to strong duality as a result. Despite this, we are able to use the cutoff prices to show that the lower bound from an arbitrary dynamic information structure coincides with the benchmark without information arrival. The remaining details are left to the appendix.

1.7.3 Case Three: Information can depend upon past (though not current) prices and is dynamic

Lastly, we consider allowing arbitrary dynamic information structures, with the restriction that information in any period not depend on that period's price, as described prior to the statement of Proposition 3. Here, the no-delay profit is not achievable. This may be expected since this is a hybrid of the previous two cases (i.e., allowing for past price dependence and multiple periods of information arrival), and since each required a different generalization of the Du mechanism:

Claim 4. *Consider a model with two periods and one buyer arriving in each period. Suppose nature can provide information dynamically (to the first buyer). Assume that $\Pi_{RSD} > \Pi^*$ and that Du's mechanism is uniquely maxmin optimal in the one-period problem. Then the seller's maxmin profit in this two-period model with arriving buyers is strictly below $(1+\delta)\Pi_{RSD}$ for any $\delta \in (0,1)$.*

While the proof of this claim is fairly involved, the information structure chosen by nature is simple. When a buyer arrives, nature provides her with the Roesler-Szentes information structure. This yields profit at most Π_{RSD} from the second buyer, and similarly from the first buyer if she expects no additional information in the second period. However, nature promises to reveal the value perfectly, in the second period, to any buyer who *would* *have* purchased in the first period without any additional information. The key technical step of the proof shows that delay always hurts the seller, and it occurs with strictly positive probability.[23] Since this requires the ability to release some payoff-relevant information to the buyer, this step is only valid if the seller utilizes random prices (as would have to be part of the optimum since $\Pi_{RSD} > \Pi^*$). Claim 4 can thus be interpreted as saying that whenever randomization is required, the one-period profit benchmark Π_{RSD} is unattainable with arriving buyers and dynamic learning.

On the other hand, we have stated Claim 4 with an extra assumption that Du's mechanism is strictly optimal. This is for technical reasons that we explain in Appendix A.2, and it may not be necessary for the conclusion. In any event, we show this assumption holds for *generic F* (see Appendix A.4.1 for details).

1.8 Informational Externalities

This section modifies the model from Section 3.1 to allow for interdependent preferences and information to be conveyed across buyers. Notice that *both* modifications must be made in order for the solution to the seller's problem to be altered. Any information generated by other buyers is simply a restriction on nature's problem, and the worst case profit could only go up by introducing constraints. And such information is meaningless unless one buyer's value influences the conditional distribution of the other buyer's value. In contrast, we will show that when both features are present, then the seller will be able to achieve a higher expected discounted profit.

We replace the independent value assumption with the other extreme, where all buyers share the same value, assuming that $v \sim F$ is drawn at the beginning of the interaction. We also assume that nature chooses a single information arrial process, $\mathcal{I}$, consisting of signal sets $(S_t)_{1 \leq t \leq T}$ and distributions $I_t : R_+ \times S^{t-1} \times P^t \to \Delta(S_t)$ that are observed by

[23] We are only able to show that for this specific information structure, total profit is strictly below $(1+\delta)\Pi_{RSD}$. Since this is generally not the worst-case information structure for every pricing strategy, we do not know how to solve for the actual maxmin profit in the model considered here.

all parties. In particular, there are two additional restrictions we are imposing on nature's choice of information arrival processes. First, a buyer that arrives at time t observes the signals observed by any buyer at time $t' \leq t$. Second, all buyers that have the opportunity to buy at time t observe the same signal realization. We return to this in Example 3.

Our first result restricts the set of relevant information arrival processes we need to consider. It turns out that for increasing price paths, it is sufficient to provide a single signal at time 1, which is observed by all subsequent buyers:

Lemma 3. *Consider the model with common values and public signals suppose buyers are short lived. Fix a weakly increasing price path $(p_1, \ldots, p_T)$ with $p_1 \leq p_2 \leq \cdots \leq p_T$. Then the worst case profit is achievable by an information structure that involves a single signal that is observed by all buyers.*

This result is the analog of Lemma 2 for this setting. The steps are similar, but the exact form of the replacement differs since buyers share the same set of information. Instead, we conclude that the worst case is achieved by in an information structure consisting of a single public signal.

Lemma 3 implies that for increasing price paths, it is sufficient to look at a single distribution of posterior valuations to find optimal price paths. We can write the seller's discounted profit from an increasing price path as:

$$\Pi^C = \min_{\tilde{F}} \sum_{t=1}^{T} p_t \delta^{t-1}(1 - \tilde{F}(p_t)), \tag{1.4}$$

where $\tilde{F}$ is some distribution of posterior expected values arising from an information structure.

Against decreasing price paths, the partitional threshold information arrival processes we used in the independent values setting can also be used here to show that Π^* is the most the seller could achieve with a declining price path. In general, however, the seller can do better with a strictly increasing price path. [24] Intuitively, since information is shared across

[24] This result was included in an earlier version of this paper, which also showed that introductory pricing is strictly optimal with two periods.

periods, using the worst-case partitional information structure against p_1 (with threshold y_1) makes it impossible to use the worst-case information structure against $p_2 > p_1$ when $p_2 < y_1$. While we conjecture that introductory pricing is in fact strictly optimal, the difficulty is to a provide meaningful upper bound for the seller's maxmin optimal profit against non-monotonic pricing strategies. This was not a concern in the independent value setting, because information structures were constructed buyer-by-buyer. On top of this, it is difficult to conjecture a tight upper bound to the profit guarantee, at least for fixed parameters.

However, we *are* able to show that there is a sequence of introductory pricing strategies that become optimal as $\delta \to 1$ when $T = \infty$. We do this by appealing to a different upper bound on the seller's profit, namely Π_{RSD}, which can be ensured if nature ignores price dependence:[25]

Theorem 5. *Consider the model with common values and public signals. Let $\Pi^C(\delta, T)$ be the seller's optimal payoff and discount factor δ and time horizon T. We have:*

$$\lim_{\delta \to 1} (1 - \delta)\Pi^C(\delta, \infty) = \Pi_{RSD} \tag{1.5}$$

which is achieved by a sequence of introductory (i.e., strictly increasing) price paths.

The proof of this theorem uses that an upper bound for the left hand side of (1.5) is obtained via a public Roesler-Szentes information structure for all buyers. We then use our (random price) Du mechanism to construct a sequence of price paths such that as $\delta \to 1$, the expression for the average per-period profit converges to the single period profit under a Du mechanism. These price paths, for uniformly distributed values, are shown in Figure 2, for $\delta = 9/10$ and $\delta = 95/100$ (fixing the initial price to be Π_{RSD}). We see that they involve the monopolist raising prices steeply at first, eventually leveling out. When the public signal is the Roesler-Szentes information structure, then the probability that the monopolist sells in every period is bounded away from 0, since the largest value in the

[25]The same result could be obtained if we replaced the left hand side of (1.5) with $\lim_{T \to \infty} \frac{1}{T} \Pi^C(1, T)$

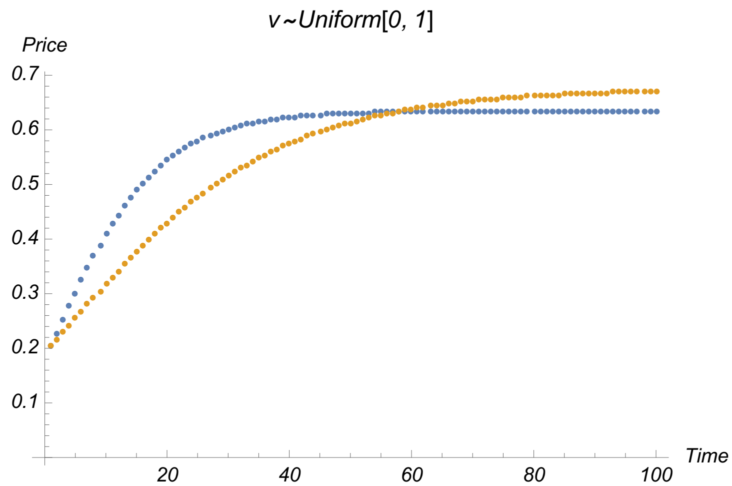

Figure 1.1: *Illustration of constructed price paths. Blue is $\delta = 9/10$, orange is $\delta = 95/100$.*

support of the constructed price distribution is strictly below the supremum of the support in the Roesler-Szentes distribution (see the Appendix for details).

The problem of optimal pricing when information is conveyed across buyers has been studied in several other papers in other Bayesian settings, such as Bose et al. (2006, 2008) (as well as papers cited therin). A major difference between this literature and our setting is that we allow buyers to delay purchase. While some papers do allow for buyers to make decisions over time, as far as we are aware, these all involve "small" buyers who have a negligible impact on the information structure. This feature does not hold in our setting, since delay by any buyer would influence information arrival. One can show that for increasing price paths, the profit levels coincide in the worst-case.

We conclude by discussing the assumption that information is public (in that all buyers observe the same signals). The following shows that without this assumption, the possibility that buyers observe imperfectly correlated signals can lower the profit guarantee:

Example 3. *Take $T = 2$, and suppose the common value is $v \in \{0, 1\}$ with $\mathbb{P}[v = 1] = 1/2$ with $\delta = 1$ (the same conclusion will hold for δ sufficiently high). Suppose the seller utilizes prices $p_1 = 2/5$ and $p_2 = 1/2 > p_1$. With short lived buyers, if the first buyer observed the second buyer's signal following any delay, the worst case information structure would involve a signal observed by* *both buyers, whose posterior expected value is supported on* $\{2/5, 1/2, 1\}$. *Under this restriction, one can show that the worst case information structure induces posterior value equal to* $2/5$ *with probability* $5/6$, *and posterior value of 1 otherwise, yielding profit to the seller of* $\frac{3}{20}$.

We show that if the second period signal need not be public, the seller does strictly worse. First, suppose no information is provided to the first period buyer in the first period, and no information is provided to the second buyer in the second period. However, in the second period, the buyer from the first period obtains a signal such that the probability of having posterior expected value 1 in the second period is $1/5$. *The first buyer is willing to delay, since purchasing in period 1 yields payoff* $1/10$, *whereas the payoff from delay is* $(1/5) \cdot (1/2) = 1/10$. *In this case, the seller's profit is* $\frac{1}{10} < \frac{3}{20}$.

This example shows that the equivalence to the short-lived benchmark of Boes et al. (2006, 2008) requires public information. Whether information is public *within the same period* is only matters for profit when delay is possible, and when buyer behavior influences information. Example 3 emphasizes distinction between public signals and outcome observations.

1.9 Our Robustness Concept

1.9.1 Seller Initial Information

Our model so far assumes that the seller has no knowledge over the information the buyer receives. In practice, however, the seller may know that the buyer has access to at least some information. For example, he may conduct an advertising campaign, and understand its informational impact very well (Johnson and Myatt (2006)). While it may be impossible or difficult for such an advertising campaign to remove all uncertainty, the seller may nevertheless know that the buyer has access to some baseline information.[26] In this section we show that this possibility does not change our conclusions.

We modify the model in Section 3.1 by assuming that in addition to having the prior

[26]Note that if the seller has complete control over what information he provides, it would be impossible to do better than the full information outcome because nature could always reveal the value.

belief F, the buyer observes some signal $s_0 \in S_0$ at time 0. The signal set S_0 as well as the conditional probabilities of s_0 given v are common knowledge between the buyer and the seller, and we denote this initial information structure by $\mathcal{H}$. We allow nature to provide information conditional on s_0 but keep all other aspects of the model identical. Equivalently, the seller seeks to be robust against all dynamic information structures in which buyer learns $\mathcal{H}$ and possibly more information in the first period.

A signal s_0 induces a posterior belief on the buyer's value, which we denote by the distribution F_{s_0}. Define G_{s_0} to be the transformed distribution of F_{s_0}, following Definition 1. The same analysis as in Section 1.4 yields the following result:

Proposition 1′. *In the one-period model where the buyer observes initial information structure $\mathcal{H}$, the seller's maxmin optimal price $p^*_{\mathcal{H}}$ is given by:*

$$p^*_{\mathcal{H}} \in \operatorname*{argmax}_p \, p(1 - \mathbb{E}_{s_0}[G_{s_0}(p)]). \tag{1.6}$$

We denote the maxmin profit in this case by $\Pi^*_{\mathcal{H}}$.

The expression (1.6) is familiar in two extreme cases: if $\mathcal{H}$ is perfectly informative, then F_{s_0} is the point-mass distribution on s_0. This means $G_{s_0}(p)$ is the indicator function for $p \geq s_0$, so that $\mathbb{E}_{s_0}[G_{s_0}(p)] = F(p)$. In contrast, if $\mathcal{H}$ is completely uninformative, we return to Equation (1.1).

For the multi-period problem, our previous proof also carries over and shows that the seller does not benefit from a longer selling horizon.

Proposition 2′. *In the multi-period model where the buyer observes initial information structure $\mathcal{H}$, the seller's maxmin profit against all dynamic information structures is $\Pi^*_{\mathcal{H}}$, irrespective of the time horizon T and the discount factor δ.*

The proofs of these results are direct adaptations of those for the model without an initial information structure. Thus we omit them from the Appendix.

1.9.2 Information versus Taste

An equivalent (albeit more abstract) formulation of our model would be to assume that each buyer observes her value perfectly, but where the value follows a *stochastic process*, with the buyer instead drawing a vector of values $(v_1, \ldots, v_T) \sim F$, where each coordinate denotes the value from purchasing the object at a given time. Nature's problem can then be thought of as choosing the stochastic process from some set of possible stochastic processes. Of course, *some* restrictions on the stochastic processes must be made to avoid a degenerate solution. However, one could instead follow Carrasco et al. (2017) and assume instead that the distribution of buyer values must have some fixed mean and support. We give an example of such a calculation in the Appendix, where we show that if the seller only knows $\mathbb{E}[v] = 1/2$ with the distribution having support [0,1], the robustly optimal price is $\frac{1}{2}(2 - \sqrt{2}) > 1/4$ and profit is $\frac{3}{2} - \sqrt{2} < 1/8$.

In this section, we argue that the restriction on how the buyer's value evolves over time (as implied by information arrival) is the significant feature that separates information arrival from taste shocks. Consider the following formulation with a single buyer that highlights the restrictions imposed by information arrival. For simplicity, we focus on the case of $T = 2$. Suppose the first period buyer's surplus from purchasing in the first period, v_1, is distributed according to $\tilde{F}$, a distribution chosen by nature to minimize the seller's profits given the pricing strategy. Suppose, in addition, $\tilde{F}$ must second order stochastic dominate the distribution F (as in our single period model). Now let the second period value be $v_2 \sim \tilde{F}(\cdot \mid v_1)$, where this random and need *only* respect the condition that $\mathbb{E}[v_2 \mid v_1] = v_1$. In particular, we do not assume the second period value changes due to information arrival. The second buyer's problem is as before.

In this setting, the proof of Proposition 2 carries over without change, since any stochastic process for the buyer's value can be induced by the dynamic information structures we construct. But actually the seller can be hurt much worse when there are arriving buyers:

Claim 6. *Consider the two period model, with one buyer arriving in each period, and suppose each buyer's initial value is distributed according to a distribution $\tilde{F}$ that second order stochastically* *dominates F. Suppose the only restriction on the first period buyer's value in the second period is that the expectation is the equal to the first period problem. Then the maxmin optimal profit for the seller when there are two buyers is* Π^*. *Hence the seller does not benefit at all from the presence of a second buyer.*

This case is extreme due to the lack of support restrictions for the second period value, although similar (but less sharp) results could be obtained under more stringent restrictions on the evolution of the buyer's value process. When nature is not restricted by information arrival, non-purchasing buyers delay because their second period value may result in negative surplus from purchasing (similar to Section 1.7.3). In particular, in this benchmark, we are not able to "push the recommendation to time 1" as we are in Lemma 2. We consider a value process such the first period buyer's purchase probability is approximately 0 as long as $\delta > 0$ against a constant, non-negative price path. But for information structures, we cannot have $\mathbb{E}[v \mid \text{Don't buy}] \leq p$ and $\mathbb{P}[\text{Don't buy}] \approx 1$ whenever $p < \mathbb{E}[v] < \infty$.[27]

For general models with taste shocks, it may not make sense to restrict the mean in the second period to be equal to the mean in the first period. The restriction was made here to relate the model where the value evolution is restricted by information arrival to one where it is not. Finding sensible restrictions on the evolution of values under taste shocks that avoids degenerate solutions (or more general conditions that would yield our same results) is left to future work.

1.10 Conclusion

In this paper, we have studied optimal monopoly pricing with dynamic information arrival while utilizing a maxmin robustness approach. With no known informational externalities, we show that the monopolist's optimal profit is what he would obtain with only a single period to sell to each buyer, and that a constant price path delivers this optimal profit. The

[27]This is seen by recalling that for a given price, the probability a buyer does not purchase is maximized by a partitional information structure which tells the buyer whether or not the value is below $F^{-1}(G(p))$.. Whenever $p < \mathbb{E}[v] < \infty$, there must be some positive probability that v is above this threshold.

inability to condition on a buyer's arrival time therefore imposes no cost on the seller (in our main model). These conclusions depend on our assumptions regarding the timing of information release, and we have illustrated how this is the case.

The question of how to design optimal mechanisms under general information structures is one that we hope will be studied in other contexts and under other modeling assumptions. One could also ask similar questions in settings where the agent's problem may not be represented by the choice of a stopping time. And settings with competition, changing values, richer population dynamics and different seller objectives all seem intriguing as well.

This paper contributes to a growing literature which employs the maxmin approach in analyzing the optimal design of mechanisms. The literature has mostly focused on static settings, although we suspect dynamic settings will receive significant attention in the future. For us, the maxmin objective is useful in two respects:

- Motivating our focus on partitional information structures, and
- Simplifying the set of relevant information structures with increasing price paths.

We hope our analysis has suggested ways that such models could be analyzed to produce new economic insights.

Chapter 2

False Positives and Transparency in Scientific Research

Any analysis that relies upon statistical inference inevitably risks arriving at an incorrect conclusion. Nevertheless, as argued in Ioannidis (2005), there are compelling reasons to believe that mistakes in published research cannot be explained by statistical error alone, and arise due to decisions in experimental design leading to bias (with bias referring to a higher probability of positive results, irrespective of hypothesis validity). A natural question arises as to how to evaluate policies designed to combat false positives and improve research quality.

Toward that end, the medical research journal *The Lancet* published a series of articles in January 2014 discussing guidelines to improve the efficiency of scientific research.[1] The prevalence of false positives was one particular focus, and partially attributed by Ioannidis et. al. (2014) to the lack of documentation requirements for experimental conduct. During any experiment, researchers make a number of decisions that could lead to bias, and not all of them will be described in the resulting publication. Examples discussed by Ioannidis et. al. (2014) include selecting samples for high-risk patients, failing to adapt significance levels

[1] The biologist Ed Wilson wrote a letter in response that, while praising the series in general, specifically lamented the lack of involvement from economists.

to the number of tests run, selective reporting of results, or insufficient training of clinical researchers.

Ioannidis et. al. (2014) suggest implementing registration requirements to combat these issues, with the understanding that scientists typically need some kind of external certification to make credible claims regarding research activity.[2] But given the varying degree of difficulty in verifying distinct activities that could lead to bias, these requirements are not a simple yes-or-no matter. Instead, implementation would occur on specific dimensions of an experiment. It may be easy for some outside authority to verify properties of certain samples. But it may be difficult to verify that research assistants were well-trained, or to distinguish a genuine need to restart an experiment from disappointment with a negative result. Answering whether a particular registration policy is beneficial requires an analysis of how the properties of experiments will change in response.

Stepping back, however, followers of the applied mechanism design literature could be skeptical that transparency requirements necessarily lead to improved research output. To see why, note that limited contractibility is the rule, not the exception, for many types of scientific research. A researcher studying whether a particular gene can lead to a particular kind of cancer may not know which pharmaceutical company would be able to use that insight in order to develop a treatment. And fundamental research is often left to universities when this research requires a time horizon much longer than private entities would be capable of providing.

In settings with limits on contractibility, several authors have cast doubt on the optimality of transparency requirements, even *without* resorting to costs of monitoring.[3] As discussed in the literature review, Prat (2005) notes that transparency requirements are eschewed

[2] Some researchers have even considered the effects of similar policies in economics and social science (Miguel et al. (2014), Coffman and Niederle (2015), Olken (2015)).

[3] Of course, if one truly thought transparency requirements contributed to false positives, one could always assert that false positives are a priori problematic and should be minimized. But as Glaeser (2006) notes, this is unlikely a natural objective. This paper adopts his view that, if false positives were widely known to be prevalent (evidenced by the number of citations of Ioannadis (2005)), observers should read results skeptically and debias accordingly.

in a variety of settings, such as corporate governance, where one may expect them to be beneficial. His model shows that agents can be more willing to act according to their private information under non-transparency in a career concerns model. Similarly, Cremer (1994) ties transparency requirements to the ability to provide high powered incentives within organizations. His model explains some documented features regarding firm boundaries.

This paper comments on the efficacy of transparency requirements in a simple framework that can shed light on the aforementioned debate in scientific research communities. In our model, a scientist (she) chooses an experiment that is characterized by a vector of research activities, and produces an observable outcome (success or failure) that is seen by a developer[4] (he). One dimension could reflect the number of samples collected; another, the number of times the experiment is repeated; or whether they are disingenuous and decide to directly alter their data to increase the probability of a positive result. The experiment imposes a cost on the scientist, but provides information as to whether or not the developer will be able to successfully develop a drug (which would yield a benefit to both players) by exerting costly effort.

This paper focuses on the question of whether the inability of the developer to observe experimental methods (due to lack of transparency) makes him better off or worse off. The main result highlights that whether transparency requirements are advantageous depends on the complementarity in costs between different kinds of research actions.

To illustrate the intuition, consider the discovery of the Higgs Boson in 2012 using the Large Hadron Collider (LHC) at CERN. Discovery actually meant that five-sigma confidence (a chance of roughly 1 in 3.5 million) had been reached.[5] In other words, the existence of the Higgs Boson was a statistical finding, inevitably short of mathematical certainty (though, depending on opinion, perhaps not by much). The discovery was a remarkably high profile

[4]We use “developer” here in order to make our story more concrete. For interpretation purposes, we can think of this person as someone intrinsically interested in the results of the scientist’s experiment, for whatever reason.

[5]Actually, three sigma confidence had already been obtained. That level of confidence is uncommon in other disciplines.

event that garnered widespread celebration, and culminated in a Nobel Prize for Englert and Higgs the following year (Wired (2015)). However, the LHC was shut down for two years after the discovery, preventing replication. Furthermore, the data from the experiment were not released for another four years, and leaving the algorithms used in the particle's discovery unexamined by the full research community.

Why was the physics community apparently unconcerned by the lack of public data and the inability to replicate the experiment? Presumably the lack of transparency was not a concern due to the difficulty in being able to bias an experiment with a five-sigma significance threshold. The five-sigma threshold is common in physics due to the feasibility of collecting large amounts of data (particularly at CERN). This does not apply to other empirical disciplines like biology. Still, it *is* often the case that a dataset's size is left to the discretion of researchers. But due to limited contractibility over experiments, scientists typically cannot be compelled to incur the costs of additional data collection directly.

This paper highlights that transparency requirements have the benefit of discouraging bias, but might also discourage scientists from undertaking costlier experiments to counteract the perception of bias. Imagine a researcher deciding to collect a large dataset (e.g., the equivalent of an experiment at CERN) or a small dataset. Large dataset experiments may be much more difficult to undertake than small dataset experiments. If the scientist is required to make all data and regression specifications public, then they could very well prefer the small dataset experiment, believing it is informative enough and finding the large dataset too costly. Without transparency requirements, however, an outsider observing a positive result would assume that some kind of biasing activity (e.g., regression fishing) took place and therefore not adapt his beliefs to the outcome of the experiment. Since the large dataset experiment is less susceptible to this kind of bias (as appeared to be the case for the Higgs Boson discovery), this loss in credibility would then induce the scientist to compensate by using the larger dataset. This compensation, in turn, can overwhelm the loss of informativeness due to biasing.

Framed in this way, a lack of transparency requirements encourage the use of large dataset experiments (such as those performed at CERN) by instead *discouraging* the use of less informative experiments that were susceptible to bias. In general, the strength of these complementarities determine whether non-transparency simply adds bias or can be compensated for more effectively along other dimensions. This result is similar to those on the optimality of "money burning" that can arise in delegation settings (for instance, as in Ambrus and Egorov (2014)). However, this conclusion is only true if the scientist's preference for developer success is sufficiently important, since otherwise non-transparency will not have this beneficial effect.

This paper contributes to the aforementioned policy debate by clarifying the connection between preferences of scientists and the merits of transparency requirements. There are two key features of the model which drive our conclusions: First, scientists care about follow-on research, and second, difficulty or costs associated with experiments influence experiment choice. Debates over transparency requirements should take these incentives into account. For instance, new researchers may be preoccupied with making tenure or developing a reputation in the discipline, while established researchers may be influencing drug companies directly and want to ensure that their research is useful. Depending on the form these career concerns take, our analysis in Section 2.4.2 suggests that it may be undesireable to apply the same transparency requirements for early-career grants versus late-career grants.

Finally, despite our focus, there is no reason the model should only be applicable to scientific research. But the justification of preferences and limits to contractibility seem appropriate for this application. And as mentioned above, transparency over differing research activities is an active policy question. While the model could certainly describe other kinds of information acquisition, we leave a more thorough analysis of these applications to future work.

We proceed as follows. We first discuss the literature and introduce the model in Section 2.2. Readers interested in the mechanism highlighted but not the general model are encouraged to skip to Section 2.3.1, which show the key forces at work. In Section

2.3.2, we consider the scientist's equilibrium behavior, and in Section 2.3.3, provide a cost function for any arbitrary experiment set such that drug developers are better off without transparency requirements. We use the model to comment on policy in Section 2.4. We proceed to consider a number of alternate specifications for the analysis in Section 2.5, and conclude in Section 2.6. Most proofs are in the Appendix.

2.1 Related Literature

The non-transparency result provided in this paper is similar to a number of others that have been derived in the literature. Results of this form can be found in Cremer (1994), Prat (2005) and Bergemann and Hege (2005). In these papers, the intuition behind the optimality of non-transparency is that it gives the principal additional commitment power which would not be credible under full transparency.[6] The results of this paper are in a similar spirit, but allow for a variety of research actions, and also provide a central role for the incentives to induce follow-on work. A more direct contrast is that the present paper obtains non-transparency as a way of burning surplus in a way that aligns the incentives of the principal and the agent. The presence of another action that is incentivized by this "money burning" is crucial for our result, and does not have a direct counterpart in these papers. Importantly, this feature relies upon the high degree of alignment between preferences of scientist and developer.

When framed as a result of money burning, our results are therefore actually closer to Angelucci (2014) or Szalay (2005), whereby the incentives are aligned when the principal takes actions which seem to harm both players. But in these papers, the distortions take other forms. Ambrus and Egorov (2014) consider cases where money burning is part of an optimal contract in a delegation setting, though our counterfactual of non-transparency does not nest in their framework.

The model itself is reminiscent of multitasking, Bayesian Persuasion (a la Kamenica and

[6]Other papers give conditions under which the principal is better off when the *agent* is not able to perfectly observe a state variable, for example, Jehiel (2014) and Ederer, Holden and Meyer (2014).

Gentzkow (2011)) and career concerns. Multitasking arises due to the variety of research actions the scientist undertakes, resembling the literature following Holmström and Milgrom (1991). This literature has showed that transparency over different dimensions can distort an agent's effort choice. Still, as these papers typically involve full observability of the agent's actions as achieving the first best, the contribution in this paper seems distinct.

The literature on communication games (and Bayesian Persuasion in particular) has flourished recently and been used for many applications. In fact, the idea of using these models to study scientific research is not in itself novel, as it is also done by Kolotilin (2015). The main difference with these settings is (1) we do not allow all information structures to be feasible, (2) we impose costs that are parameterized with the information structures, and (3) we consider a case where the sender can only commit to *part* of the information structure. This third point is perhaps the most significant departure. Hoffmann, Inderst and Ottaviani (2014) are also interested in the lack of commitment, but view it as arising from a disclosure problem. Several papers have consider cases where distortions of information can take particular forms, such as fraud as in Lacetera and Zirulia (2008), or selective disclosure as in Henry (2009) and Felgenhauer and Schulte (2014). Other papers have studied information aquisition and communication in cheap talk models; see Argenziano, Severinov, and Squintani (2014), as well as Pei (2014).

This introduction of signal distortion is reminiscent of the career concerns literature (Holmström (1999) and Dewatripont, Jewitt and Tirole (1999)). While we accomodate these preferences for the scientist, our model provides a novel channel for preferences over informational content (as opposed to just the outsider's posterior). We distinguish these incentives from one another. Incentives for information acquisition is represented by the *convexity* of the scientist's expected payoffs as a function of the developer's beliefs. On the other hand, the marginal benefit from distorting can be most clearly seen by studying the *slope* of the expected payoff conditional on the state (i.e. the truth of the hypothesis).

Lastly, this paper relates to the literature on academic publication (as in Aghion, Dewatripont and Stein (2005), for example). Azoulay, Bonatti and Krieger (2015) empirically study the effect of a retraction on a scientist's reputation, documenting that retractions lead to a drop in citations consistent with reputation loss. Andrews and Kasy (2017) provide methods for determining when publication bias may be important, and suggest a way to debias taking this into account. Yoder (2016) develops a principal-agent model with transfers, similarly motivated by Bayesian Persuasion, to describe the optimal incentives for research institutions. His conclusion that negative results should be rewarded is consistent with this paper, since doing so can align incentives.

Several papers have cautioned against associating false positives with problems in scientific conduct. Glaeser (2006) studies the incentives behind false positives, and argues that eliminating them may be socially harmful. His reasons for this do not directly relate to our overcompensation effect, instead focusing hypothesis choice. Kiri, Lacetera and Zirulia (2015) consider the incentives for fact-checking, arguing that failure to observe mistakes would suggest the lack of verification activities. Furukawa (2017) develops a vote-counting model to show that publication bias may arise naturally when results are only coarsely interpreted by practitioners, and suggests a way to correct for it. Hopefully these ideas will be useful in the design of research guidelines, and will further improve our understanding of how to effectively structure scientific endeavors.

2.2 Model

A scientist is endowed with a hypothesis whose validity is given by $\theta \in \{T, F\}$, drawn by nature. The scientist is able to conduct an experiment, the results of which are of interest to a drug developer.[7] All players share a common prior on θ, with $\mathbb{P}[\theta = T] = p_0$. We will always take p_0 to be interior, and we will think of T as being the "good" state, and F as the "bad" state. The hypothesis could be whether a particular gene is associated to a specific disease, or whether a certain particle can safely target a specific biological pathway.

The scientist chooses an experiment from some set of possible experiments. We assume

[7]The scientist can be thought of as a sender, or an agent, and the developer as being a receiver, or a principal.

that the set of experiments are parameterized by an n-tuple $(a_1, \ldots, a_n)$, with each $a_i \in [\underline{a}_i, \overline{a}_i] \subset \mathbb{R}$. yielding an experiment $\mathcal{I}(a_1, \ldots, a_n) : \Theta \to \Delta\{0,1\}$ at cost $c_S(a_1, \ldots, a_n)$. For now, we only assume that the cost is increasing in each coordinate, and we highlight that we will *not* assume that higher actions yield more informative experiments. As mentioned in the introduction, each dimension is meant to capture a different kind of research activity; collecting data, p-hacking activity, selecting samples, and so on.

A restriction of the above setting is that the experiment outcome can only take one of two values, so that the developer only observes experiment produces an outcome $y \in \{0,1\}$ according to a distribution that depends on θ and $a \in A = \prod_{i=1}^{n} [\underline{a}_i, \overline{a}_i]$. This outcome $y \in \{0,1\}$ is observable to the developer, and we will refer to the event $y = 1$ as a "positive result," and $y = 0$ as a "negative result." A *false positive* occurs when $y = 1$ and $\theta = F$. Define:

$$h_\theta(a) := \mathbb{P}[y = 1 \mid \theta, a],$$

and assume this function is continuous[8] with bounded derivatives. In order to interpret positive results as evidence for hypothesis validity, we assume that $h_T(a) > h_F(a)$, for all a.

The developer does not necessarily observe the entire profile of research activities a. Instead, we assume that there is some third party (e.g., a journal or a funding agency) that can dictate what the developer observes. Formally, the third party chooses a set of indices, $M \subset \{1, \ldots, n\}$ with the interpretation that the developer will observe all a_i such that $i \in M$, in addition to the outcome y itself. The interpretation is that the third party has the ability to make various dimensions, such as the number of specifications run, observable to the developer. In principle, one could imagine that there are costs associated with different choices of M, but since the implication of those costs is straightforward we do not model them explicitly. However, one could imagine that in practice these costs are prohibitive for the scientist. We are interested in how the developer's payoff changes with different choices of M.

[8]Continuity ensures the existence of a pure strategy equin order to ensure the existence of a pure strategy equilibrium with unobserved coordinates.

After observing the action choice and y, the developer updates his prior from p_0 to $\hat{p}(y)$. Since the developer's beliefs will depend on equilibrium behavior, we will typically suppress the dependence of the posterior $\hat{p}(y)$ on the experiment choice of the scientist, but occasionally we will make this dependence explicit when the choice is known to be $a = (a_1, \ldots, a_n)$ by denoting the posterior $\hat{p}_a(y)$.

The developer chooses a level of effort $e \in [0,1]$ at cost $c_R(e)$, where $c_R(e)$ is an increasing and convex function. The choice of e determines the realization of a random variable $x \in \{0,1\}$, the distribution of which is given by:

$$\mathbb{P}[x = 1 \mid \theta = T, e] = e, \quad \mathbb{P}[x = 1 \mid \theta = F] = 0.$$

We think of $x = 1$ as the event that a drug is developed, and $x = 0$ as the event that it is not. For example, the drug could attempt to cure the disease by deactivating the gene studied in the scientist's experiment. Increasing effort in drug development makes it more likely to succeed, but only in the event where the gene is in fact associated with disease incidence—that is, if the scientist's hypothesis is true.

If the drug is developed, the developer obtains a payoff of $b > 0$, and the scientist obtains a payoff of λ. As mentioned, for the baseline model we imagine there is limited contractibility between scientist and developer, so we take these payoffs as being exogenous. In that case, λ might reflect profit, whereas b may reflect prestige or pride associated with having cured a disease. We also suppose the scientist also receives a benefit of $g(\hat{p})$ when the public belief is $\hat{p}$ at the end of the game, for some increasing function g. We refer to this as payoffs coming from career concerns; insofar as scientists may (be thought to) have hypothesis validity correlated across time, then they may place some premium on having future drug developers believe that their past hypotheses were true. We use this interpretation to distinguish our specification from other papers and to discuss policy recommendations. In an appendix, we discuss a version of the model where the scientist also has preferences over y itself.

Hence final payoffs for the scientist are

$$\lambda \cdot x + g(\hat{p}) - c_S(a),$$

whereas for the developer, they are

$$b \cdot x - c_D(e).$$

We assume that the parameters are such that the developer's optimal effort choice always involves $e < 1$. [9]

2.2.1 Measuring Uncertainty

In this section we place some restrictions on how we parameterize information structures. We also introduce the following terminology:

Definition 2. *A dimension a_i is biasing if $h_T(a_i, a_{-i})$ and $h_F(a_i, a_{-i})$ are both increasing in a_i, for all a_{-i}. A dimension a_i is informative if $h_T(a_i, a_{-i})$ is increasing in a_i and $h_F(a_i, a_{-i})$ is decreasing in a_i, for all a_{-i}.*

For the main results of the paper, we will assume that experiments are Blackwell ordered along each dimension (increasing along informative dimensions and decreasing along biasing dimensions). This assumption is stronger than what is necessary, although without it, we would need stronger assumptions on preferences in order to ensure that actions could be classified as biasing or informative. For example, our results would continue to hold if we assumed quadratic developer effort costs and considered experiments ordered by the reduction in posterior variance (which is implied by Blackwell ordering of experiments but not conversely). We adopt this specification for analysis of Examples 4 and 6 below.

Note that the developer's expected payoff, given effort e and belief p, is given by:

$$\pi_R(\hat{p}) := \max_e b\hat{p}e - c_R(e).$$

[9] A sufficient condition that ensures this, for example, would be that $c_R(e) = \frac{k}{2}e^2$ and $b < k$.

We adopt the terminology of Ely, Frankel and Kamenica (2015) to call a measure of uncertainty to be a strictly concave function that is 0 at degenerate beliefs. Define the *developer's measure of uncertainty* from an experiment $\mathcal{I}$ to be $p_0\pi_R(1) - \mathbb{E}_{y\sim\mathcal{I}}[\pi_R(\hat{p}(y))]$.

Proposition 4. *The payoff gain of the developer from observing the scientist's experiment is a measure of uncertainty.*

It is immediate from the definition that the developer's payoff is higher for experiments that are more informative according to his measure of uncertainty. The proposition, together with our ordering on experiments, implies that the optimal *scientist* experiment for the *developer* would involve the maximal action on informative dimensions and the minimal action on biasing dimensions.

2.2.2 Examples of Information Acquisition Technologies

To illustrate the analysis, we provide two stories behind information acquisition technologies which satisfy the assumptions of the model. These are meant to demonstrate a tighter link between what the model captures and the kinds of scientist behavior practitioners tend to be concerned about.

Example 4 (p-hacking)**.** *First consider the case of $c_R = \frac{k}{2}e^2$, so that we can assume experiments are ordered by posterior variance (as opposed to Blackwell ordered). Suppose the scientist chooses a number of times, $a_2 \in [3, 15]$ to run an (iid) experiment and a precision a_1, where:*

$$\mathbb{P}[\textit{Success on an experiment} \mid \theta = T] = \mathbb{P}[\textit{Failure on an experiment} \mid \theta = F] = a_1.$$

The resulting informativeness (as measured by the posterior variance) is plotted for $a_1 = 2/3$ and $p_0 = 1/2$ in Figure 2.1. While it increases at first, eventually it decreases, with the posterior variance approaching 0. We can therefore think of one dimension as being the amount of p-hacking and the other to be the informativeness of the underlying experiment (which are biasing and informative, respectively, on some appropriate range).

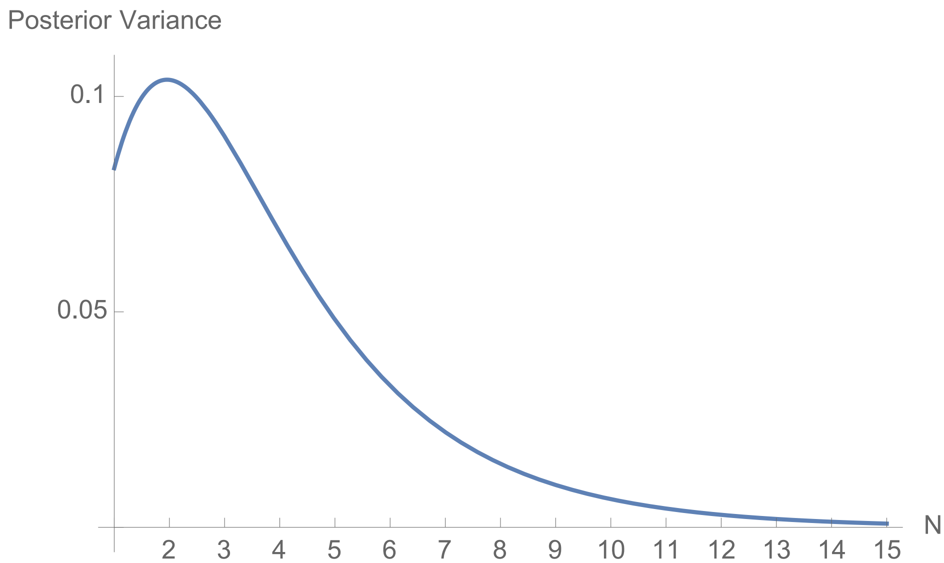

Figure 2.1: *Posterior Variance for the p-hacking specification, where N is the number of trials.*

Example 5 (Lying or Fudging Data)**.** *Again suppose that a_1 parameterizes the underlying experiment as in the p-hacking example. However, now suppose that a_2 is the probability that the scientist changes a result of $y = 0$ to a result of $y = 1$—for example, as a result of either direct falsification or altering the data. This is a Blackwell garbling of the underlying signal technology, and hence decreases the informativeness for any informativeness measure, not just the relevant one for the developer (i.e. the posterior variance).*

Example 6 (Quality Decisions)**.** *Now suppose that a_1 is investment in research equipment, and that a_2 measures the quality of lab technicians or research assistants (supposing it is as easy to hire good research assistants as bad ones). Suppose the resulting information structure is*

$$\mathbb{P}_{\mathcal{I}(a_1,a_2)}[y = 1 \mid T] = (4/5)a_1 + (1/5)a_2, \quad \mathbb{P}_{\mathcal{I}(a_1,a_2)}[y = 1 \mid F] = (1/5)a_1 + a_2,$$

where $a_2 \in [\underline{a}, \overline{a}]$ reflects employee quality and $a_1 \in [\underline{q}, \overline{q}]$ reflects equipment quality, with parameters chosen so that probabilities are within $[0, 1]$ and $\mathbb{P}_{\mathcal{I}(a_1,a_2)}[y = 1 \mid T] > \mathbb{P}_{\mathcal{I}(a_1,a_2)}[y = 1 \mid F]$.

Of course, our model would also accomodate hybrids of these technologies, and most of the results turns out to not depend on the particular functional forms of the set of experiment space. The underlying point is that the model can accommodote natural qualitative features of experimentation, treating different kinds of research activities as observationally distinct.

2.3 Main Results

2.3.1 A Simple Example

In this section, we illustrate our construction and walk through the intuition in a numerical example. We take $p_0 = 1/4$, and specialize to the quality example of Example 6. Take we take $c_R(e) = \frac{1}{2}e^2$ and $b = \lambda = 1$. Suppose the experiments the scientist has access to are as follows, for $a_1, a_2 \in \{0, 1\}$:

$$\mathbb{P}_{\mathcal{I}(0,a_2)}[y = 1 \mid T] = 2/5 + (1/10)a_2, \qquad \mathbb{P}_{\mathcal{I}(0,a_2)}[y = 1 \mid F] = a_2/6,$$
$$\mathbb{P}_{\mathcal{I}(1,a_2)}[y = 1 \mid T] = 4/5, \qquad \mathbb{P}_{\mathcal{I}(1,a_2)}[y = 1 \mid F] = 0,$$

where $\mathbb{P}_e$ denotes the probability measure when experiment e is chosen. One can check that as a_2 increases, the experiment $\mathcal{I}(0, a_2)$ becomes less informative in the sense of Blackwell. Further note that a_2 does not affect the informativeness of the experiment $\mathcal{I}(a_1, a_2)$.

An interpretation of this parameterization comes from Example 6, where a_1 may reflect equipment quality and a_2 may reflect training or quality of research assistants. The idea is that experiments with "low quality equipment" are both less informative, and susceptible to bias depending on the research assistants. In contrast, "high quality equipment" gives an experiment which does not require the help of research assistant, and is also more informative no matter what. The important feature of this example is that while these actions might make an experiment susceptible to false positives, a positive result is already very likely if the equipment is of high quality, *provided* the hypothesis is true. We suppose experiments with low quality equipment are costless, but high quality equipment is costly. Research assistant quality, in contrast, is costless.

Suppose the developer sees the complete experiment chosen by the scientist—both equipment quality a_1 and research assistant quality a_2. For instance, suppose the funding agency requires documentation of equipment and can require research assistants to satisfy certain prerequisites. Discreteness is helpful in that we can compute payoffs experiment-by-experiment: For $i \in \{S, R\}$,

$$\pi_i(0,0) = 1/8 \qquad \pi_i(0,1) = 1/12$$
$$\pi_i(1,0) \approx 1/5 \qquad \pi_i(1,1) \approx 1/5,$$

where $\pi_S(a_1, a_2)$ and $\pi_R(a_1, a_2)$ are the benefits to the scientist and developer, respectively, from an experiment of type (a_1, a_2).

Now suppose that the developer can observe equipment quality a_1, but not research assistant quality a_2, leaving research assistant quality unverified. In this case, if the scientist picks $a_1 = 0$, then $a_2 = 0$ will not be chosen when the choice of a_2 is costless. To see this, simply note that the scientist's expected payoff is always higher following a signal of $y = 1$ than a signal of $y = 0$. On the other hand, the developer cannot distinguish $a_2 = 0$ and $a_2 = 1$, and so in equilibrium his belief will not change with the choice of a_2. Since higher a_2 generates a higher probability of $y = 1$ when $\theta = T$, the scientist opts for lower quality research assistants.

Does this mean that the experiment the scientist chooses is less informative when d is unobserved? Not necessarily, due to costs. If the cost of high quality equipment is *either* sufficiently small or sufficiently large, the developer does best if both a_1 and a_2 are observed. If the costs of high quality equipment are small, then the scientst would choose this experiment in either regime and hence bias due to research assistants would not be a concern. On the other hand, if the cost of high quality are sufficiently large, then the scientist will never choose these experiments and hence all non-transparency does is induce added bias. For an intermediate range of costs[10], however, the scientist switches choice of a_1 across the regimes. Under observability, low quality equipment is

The example highlights three main messages which are useful for understanding our

[10]Specifically, costs between 3/40 and 7/60 (approximately).

main results. First, the scientist exhibits a *strict preference for more informative experiments.*[11] Specifically, the above calculations verified that the scientist would always prefer a more informative experiment if the extra information were free. Our analysis in Section 2.3.2 shows that this is a general feature for a wide variety of settings. Notice, however, that the scientist does not gain anything if the developer's beliefs are unduly optimistic when $\theta = F$, and in fact strictly prefers that the developer know the state when $\theta = T$. Furthermore, the more informative the experiment, the higher beliefs are (in expectation) when $\theta = T$. We will show how the state dependent nature of the scientist's payoff as a function of beliefs translates into an incentive for information acquisition, which will allow us to distinguish these incentives from the incentives for distortion.

Second, despite liking informativeness, when a_2 is unobserved the scientist *loses credibility for scrupulousness*. Since beliefs do not respond to the choice of distortion, and since the posterior rises if and only if $y = 1$, increasing a_2 increases the probability that $\hat{p}$ is higher when $\theta = T$. Hence even though the scientist would like to set $a_2 = 0$, she cannot commit to doing so when unobserved, since the developer will realize that it is profitable for her to deviate to $a_2 = 1$. This is also studied in Section 2.3.2.

Third, despite the loss of credibility for scrupulousness, the scientist *compensates by exerting costly effort* which ultimately increases the informativeness of the chosen experiment. Since the scientist has a preference for more information, the only reason a scientist would choose a less informative experiment over a more informative one would be on account of costs. So by making it impossible to commit to $a_2 = 0$, the scientist is induced to take a costly action which proves her scrupulousness. Here, this takes the form of choosing an even more informative experiment. In other words, the scientist needs to exert more costly effort (in this case, acquire higher quality equipment) in order to prove that the experiment is actually informative.

We are able to use the intuition gained from this example to demonstrate the optimality of

[11] This does not follow immediately from Blackwell's Theorem, since the scientist is not the decision maker who uses the information. See Kim (1995) for a discussion of the use of Blackwell informativeness in a principal-agent model with moral hazard.

partial transparency without having to resort to particular specifications of the informational environment. Despite this, the example is still somewhat restrictive. It rules out, for instance, higher a_2 having any influence on the experiment with higher a_1. It is also ad hoc in that it rules out multiplicity and the possibility of mixed strategies. Under the generality of the main model, these may not be ensured, and in these cases, the intuition from above cannot be applied directly. Our main results show that these can be ignored under certain specifications for cost functions. The generality of the full model is also useful in explaining why our result could not occur in other common benchmarks studied in the literature (that is, without both costly communication and limited commitment). While this may not be obvious from this specialized setting, the comparison is more direct in our general model.

2.3.2 Influence of Transparency on Experiment Choice

This section describes the how observability of a_i influences the experiment choice of behavior. We demonstrate how to adopt the belief-based approach to this setting. Gentzkow and Kamenica (2014) remarked that this is less straightforward when there are costs, since the payoff may depend on the signal structure outside of the beliefs they induce. Nevertheless, we are able to use this for our specification.

The scientist takes the effort choice of the developer as given, and the expected payoff from an experiment is a function of the posterior belief of the developer following y, denoted by $\hat{p}(y)$. Payoffs are therefore:

$$p_0 \mathbb{E}[\lambda e(\hat{p}(y)) \mid \theta = T] + \mathbb{E}[g(\hat{p}(y))] - c_S(a) \tag{2.1}$$

The following Lemma rewrites the payoffs of the scientist without conditioning on the state. It shows that we can changes the prior belief p_0 in (2.1) into a posterior belief by removing the conditioning event $\theta = T$:

Lemma 4. *In any pure strategy equilibrium where the experiment is correctly inferred as $\mathcal{I}(a)$, the* *scientist's payoffs can be written as:*

$$\mathbb{E}_{y\sim\mathcal{I}(a)}[\lambda\hat{p}(y)e(\hat{p}(y)) + g(\hat{p}(y))] - c_S(a) \tag{2.2}$$

This lemma is similar in spirit to many arguments that have utilized the belief-based approach in the persuasion literature (and is not particularly complicated), though we are not aware of (2.2)[12] having been explicitly stated or utilized directly. While Kamenica and Gentzkow (2011) do allow for state dependence on the Sender's utility function, their characterization of the Sender's value function does not require them to explicitly state Receiver's preferences, meaning this lemma would have limited use for their exercise. However, see Section 4 of their paper for a discussion of how preference alignment influences the solution.

The Lemma can intuitively be thought of as "flipping the order of integration"; rather than first taking an expectation over the signal realization y and then the state realization θ, it takes the expectation of the state *first*, and signals *second*. That said, the result relies upon the developer and scientist updating their beliefs in the same way in response to the experimental outcome. Equation (2.2) also clarifies that λ generates preference for information acquisition. It is well understood from the persuasion litearture that convexity of payoffs in beliefs can generate incentives for information acquisition, but it is hard to see how to apply this intuition directly from looking at equation (2.1). In contrast, the Lemma implies shows that the term driving the scientist's added incentive for information acquisition from the developer's actions is given by the term $\hat{p}e(\hat{p})$, meaning there are additional incentives for information acquisition whenever this term is concave.

The lemma allows us to adopt the belief-based approach in cases where the experiment choice is observed by the developer, but does not do so under different transparency regimes. In order to accomplish that task, we need to ensure that the experiment choice is *deterministic* and *unique*. If the scientist were randomizing actions, then this would imply that the scientist's beliefs following any given outcome would *not* coincide with the developer's,

[12] The appendix uses a more general identity, which is also novel to our knowledge.

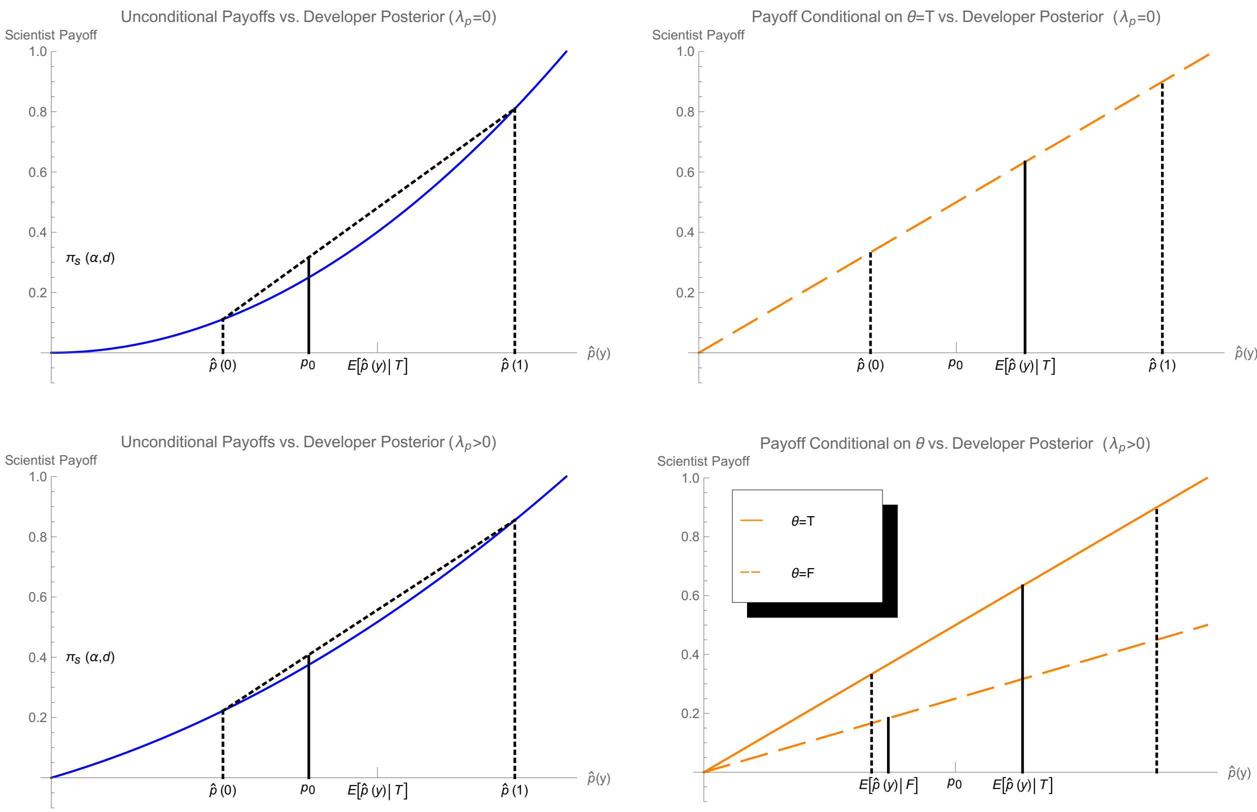

Figure 2.2: *Graphical illustration of Lemma 4. All graphs express the scientist's payoff as a function of the developer's belief. The top row considers the model where* $\lambda = 0$*, and the bottom row considers the model when* $\lambda > 0$*, where were normalize the ex ante payoff when* $\hat{p} = 1$ *to 1 for both cases. The left column displays payoff of the scientist from the ex-ante perspective, taking an expectation over the state. The right column writes the scientist's payoff conditional on each state.*

and hence we could not use a single $\hat{p}(y)$.[13] In principle uniqueness could be dispensed of we were willing to make a selection argument (as in Kamenica-Gentzkow (2011) or Lipnowski-Ravid (2017)), although given our focus on developer (receiver)-optimality it is not clear that these selection arguments would be desireable. However, using the following lemma, it follows that we can indeed adopt the belief-based approach in this setting:

Lemma 5. *Suppose that $c(a_M, a_{-M})$ is weakly convex in a_{-M} for all a_M, and furthermore, that $h_T(a_M, a_{-M}), h_F(a_M, a_{-M})$ are weakly concave in a_{-M} for all a_M.*

(1) A Perfect Bayesian equilibrium in pure strategies exists when a_{-M} is unobserved. In this equilibrium, if a_i is interior and unobserved, we have:

$$\frac{\partial c_S(a)}{\partial a_i} = \lambda \cdot (e(\hat{p}(1)) - e(\hat{p}(0)))p_0 \frac{\partial h_T}{\partial a_i} + (g(\hat{p}(1)) - g(\hat{p}(0))) \left(p_0 \frac{\partial h_T}{\partial a_i} + (1-p_0)\frac{\partial h_F}{\partial a_i} \right). \tag{2.3}$$

If a_i is observed, the first order condition corresponds to the derivative of (2.2) with respect to a_i.

(2) If either convexity of c or concavity of h_T is strict, there is no equilibrium in mixed strategies.

While the proof of this lemma is mostly straightforward, one assumption that turns out to be important is that $h_T(a) > h_F(a)$ for all a. This ensures that $\mathbb{P}[y=1]$ and $\mathbb{P}[y=0]$ are both interior, meaning that the receiver always places positive weight on observing any given signal. This is in contrast to other papers in the persuasion ltierature, which often do not restrict the set of signals that a receiver may observe a priori.

The lemma indicates that the scientist's equilibrium behavior can be thought of as equating marginal benefits and marginal costs along each dimension. While the former simply depends on $c_S(a)$, the marginal benefit is the sum of the gain due to the career concerns term as well as due to the change in the developer's effort. While such a condition is intuitive, the complication is that the marginal benefit depends on beliefs, which are endogenous. Additionally, as mentioned above, the intuition requires the validity of the

[13] In Section 2.5.1, we show that mixed strategies will generally arise with the presence of certain kinds of private information for the scientist.

belief-based approach.[14]

Another implication of the lemma is a subtle (and potentially empirically relevant) difference between the scientist's career concerns and investment in developer outcome. Lemma 4 implies that under observable behavior, the case of $\lambda = 0$ and $g(p) = \beta p e(p)$ results in an identical preference over experiments and hence identical experiment choice compared to $\lambda = \beta$ and $g(p) = 0$. However, this does not hold when dimension a_i is unobservable; in that case, the false positive rate $h_F(a_M, a_{-M})$ *does* matter for the career concerns payoff, but *does not* matter for the developer outcome payoff. We state this as follows:

Corollary 1. *When $g(p) = 0$, the false positive rate does not influence the experiment choice when action a_i is observed.*

Finally, we note that the lemma illustrates the difference between the first order condition as transparency changes. When an action changes from being unobserved to being observed (by the developer), a term equal to:

$$\mathbb{E}\left[\left(\lambda \hat{p}(y) e'(\hat{p}(y)) + g'(\hat{p}(y))\right) \frac{\partial \hat{p}(y)}{\partial a_i}\right] \tag{2.4}$$

is added to the right hand side of (2.3). This term can be positive or negative, even when information structures are monotonic in a_i in the Blackwell order; in that case $\frac{\partial \hat{p}(1)}{\partial a_i}$ and $\frac{\partial \hat{p}(0)}{\partial a_i}$ will have opposite signs, and hence (2.4) would be convex sum of a negative term and a positive term. Since we can find Blackwell ordered information structures that also hold $\hat{p}(1)$ and $\hat{p}(0)$ constant, it follows this additional term can either be positive or negative.

Corollary 2. *A sufficient condition to ensure that the scientist's payoffs is $c_R'''(e) \geq 0$, or if $c_R(e) = e^n/k$ for any $n > 1$ and $g(p)$ is not too concave. If action i is distortive, then $a_i = \underline{a}_i$ in equilibrium whenever it is observed, provided $g(p)$ is not too concave.*

[14] Admittedly this specification rules out the discrete example from Section 2.3.1, although we find it easier to parameterize the experiments more generally by resorting to the specification of the main model. Furthermore, the example is useful for the purposes of intuition in our general specification.

To summarize: the incentives for information acquisition arise due to the state dependence of the scientist's payoff as a function of the developer's posterior, since the convexity of this line is what generates incentives for information acquisition. But the loss of credibility occurs due to the positive slope of the scientist's payoff conditional on θ—that is, because the payoff is still higher when the developer's belief is higher. By devloping this model and comparing the influence of λ_p and λ_x, we have also shown why the forces highlighted are distinct from others that have been proposed, most notably in the career concerns literature.

2.3.3 Optimal Transparency

Having characterized scientist behavior as a function of the transparency requirements, we are now in a position to present results on the optimality of different transparency regimes. First, we point out that greater transparency is always better for the scientist; if an action a arises in the equilibrium where only some subset of the coordinates are observable, then the scientist could always guarantee this outcome when a is totally observable, and could in fact achieve a higher payoff potentially via some other action. This argument proves that:

Proposition 5. *Full transparency always achieves the scientist-optimal payoff.*

Our main result is that, under the payoff specifications of this setting, there generally exist payoff specifications such that full transparency is not optimal.

Theorem 7. *Suppose $c_R'''(e) \leq 0$ or $c_R = e^n/k$, $\lambda > 0$ and $g(p)$ is not too concave. Let J denote the indices that are informative and $N \backslash J$ denote the indices that are distortive, and let $K \subset N \backslash J$. Then there exists a cost function (increasing in all coordinates) such that the developer does better when a_i is not observed for all $i \in K$.*

The theorem is proved by constructing cost functions inspired by the example in Section 2.3.1. We construct cost functions with the property that distortion is "cheap" for uninformative experiments but "expensive" for informative ones. The conditions of the theorem ensure that the scientist does obtain a higher payoff from conducting experiments that are more informative. Because of this, it is costly for them to be perceived as adding bias. The cost functions constructed have the property that higher actions along the informative dimensions make biasing more costly. For example, it is hard to run an experiment many times if collecting a new data set each time is difficult. While other incentives may be present, we view this result as adding an important, subtle caution to the debate on transparency requirements.

2.4 Policy Analysis

2.4.1 Costless Communication

From the theoretical perspective, the main novelty of the model is in its use of communication costs as well as limited commitment by the sender (scientist). If any experiment is feasible and costs are not present, then the model reduces to Bayesian Persuasion (as in Kamenica and Gentzkow (2011)) when a is observable,[15] and reduces to cheap talk (as in Crawford and Sobel (1982) or Lipnowski and Ravid (2017)) in the case where a is not observable at all.[16] In this sense, the model provides an "intermediate commitment" benchmark.

In fact, we can show that without costs, the developer's payoff is increasing in the level of scientist commitment:

Proposition 6. *If the scientist's cost for all experiments are known to be zero, then making any biasing dimension unboservable strictly lowers the developer's payoffs.*

2.4.2 Career Concerns versus Follow-on Interest

Recall the earlier point (highlighted in Corollary 1) that even when career concerns replicate the preferences over experiments as interest in follow-on outcomes, there may still be differences in how the scientist responds to transparency changes. We use this insight to generate testable predictions of the model, in terms of the comparative statics of how

[15] Ichihashi (2017) studies the case of Bayesian Persuasion when the sender's choice set can be limited.

[16] Technically speaking, Kamenica and Gentzkow (2011) and Lipnowski and Ravid (2017) do not restrict to $y \in \{0, 1\}$, although it follows from their results in this setting that the sender would not benefit from a richer signal space since the state is binary.

researchers with the same access to experiments but different weight placed on immediate outcomes (as opposed to their career outcomes) would respond to transparency changes differently.

Our first observation is that, unless the payoffs derived from career concerns are strictly convex, there is no benefit to making biasing actions unobservable. This is the case since the developer has no incentive for information acquisition but does have incentives for adding bias the case; we immediately obtain.

Proposition 7. *Let $\lambda = 0$ and $g''(p) \leq 0$. Then for any cost function, full transparency is developer-optimal.*

Whether career concerns incentives should be convex or concave in general seems difficult to answer a priori. Risk aversion over long-term career outcomes would suggest concavity is appropriate, though high-power incentives for "superstar" researchers would generate convexity.

To highlight the differences between career concerns (given either convexity or concavity) and follow-on interest, we focus on the case where biasing actions do not increase the true positive rate, but only increase the false positive rate. In this case, it similarly follows immediately from Corollary 1 that:

Proposition 8. *Suppose $\frac{\partial h_T}{\partial a_i} = 0$ for any biasing action a_i and $g''(p) \leq 0$. Then full-transparency is developer optimal.*

These comparisons are significant since it is reasonable to assert that the significance of follow-on research versus the beliefs in the validity of hypotheses would vary according to the stage of the career of the scientist. For younger researchers, follow-on research may be less significant than the belief in the validity of the hypothesis. For older researchers, legacy may be more significant, in which case the importance of generating follow-on research would matter more. In this sense, the model suggests that transparency requirements would have different impacts across these different researchers. Policymakers may consider, for example, imposing transparency requirements on early stage researchers, but not late stage researcher, if they believed that most scientists were risk-averse over long term outcomes.

2.5 Extensions

2.5.1 Private Information on Distortability

So far, the set of experiments available to the researcher has been taken to be common knowledge. This is a sensible assumption if, for example, the set of possible experiments is well-understood and could be characterized in advance. On the other hand, if the scientist has specialized knowledge about the experiment in the first place, then one may also be interested in what would happen in case this assumption were relaxed.

In this section, we illustrate that the presence of this kind of private information may result in the scientist's equilibrium behavior involving mixed strategies. In order to minimize notation, we demonstrate this in the context of the example from Section 2.3.1, rather than the general model, and assert that similar intuition applies for that setting as well. Recall that in this setting, the scientist can choose a perfectly informative experiment at a cost, or can choose an imperfectly informative experiment for free (but is suceptible to bias).

Suppose instead that the scientist is only able to add positive bias to experiment with $a_1 = 0$ (low quality equipment) with probability t. This situation is illustrated in Figure 2.3, focusing again on the case with costless choice of a_2 (research assistant quality, the biasing dimension). There are two features which distinguish this version from the previous analysis. First, under non-transparency, the developer's beliefs need not form a martingale from the perspective of the scientist, when the experiment $a_1 = 0$ is chosen—that is, for the scientist, $\mathbb{E}[\hat{p}(y)] \neq p_0$. To see this, note that with probability t, the scientist chooses $a_2 = 1$ when choosing $a_1 = 0$, and with probability $1 - t$ the scientist is forced to set $a_1 = 0$. Hence the probability that $y = 1$ is larger for the scientist who is allowed to bias than it is for the scientist who cannot. But the developer cannot observe whether the scientist is able to bias, and his beliefs are a martingale from his perspective. The result is that if the scientist can

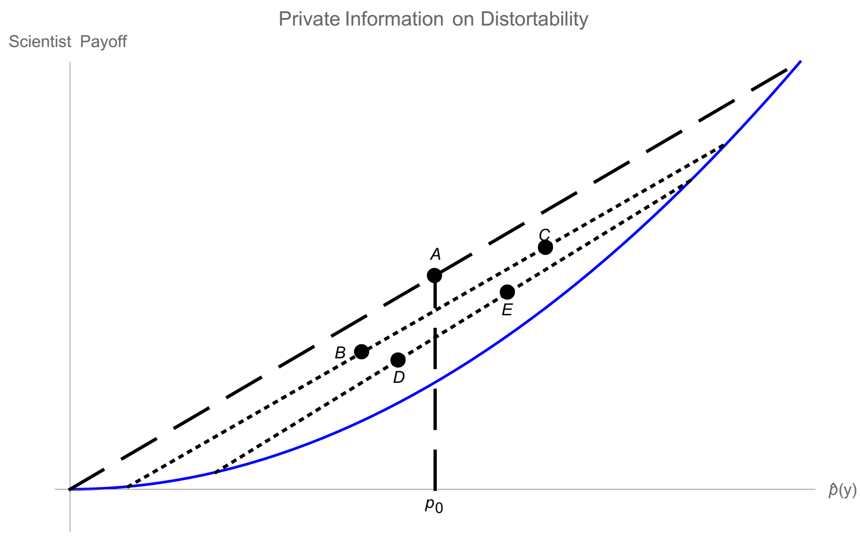

Figure 2.3: *Graphical explanation for mixed strategies in Section 2.5.1. The points C and E represent expected payoffs if distortions are feasible (hence to the right of the prior), and points B and D represent expected payoffs if distortions are infeasible (hence to the left of the prior). If the cost of the fully informative experiment moves the payoff from A to a point lower than C but higher than E, then the equilibrium behavior will involve mixed stratgies.*

pick $a_2 > 0$, then $\mathbb{E}[\hat{p}(y)] > p_0$, but if the scientist cannot, then $\mathbb{E}[\hat{p}(y)] < p_0$. [17]

To see the second difference, notice that when considering the informativeness of an experiment, we were able to treat the experiment choice as given and then ask which level of distortions would be picked. In this case, however, the informativeness of the signal $a_1 = 0$ will depend on whether a player who is able to distort would prefer to choose $a_1 = 1$ or $a_1 = 0$. If the scientist chooses the proposal $a_1 = 0$ when distortions are available, then the experiment is less informative than it would be if she were to choose $a_1 = 1$ when distortions are available. When $t = 1$, these two cases are the same, but they are not in general.

This observation explains why we may have difficulty with ensuring the existence of a pure strategy equilibrium. Indeed, for certain values of t, the scientist who can choose $a_1 = 0$ and bias might mix between $a_2 = 1$ and $a_1 = 0 = 1 - a_2$. Indeed, it may be the case that, without mixing, if the developer thought the scientist would always choose $a_1 = 1$

[17] We again emphasize again that $\hat{p}(y)$ is the probability the developer, and not the scientist, assigns to the event that $\theta = T$.

if available, the scientist would prefer to choose $a_1 = 0$ and set $a_2 = 1$, whereas if the developer thought the scientist would only choose $a_1 = 1$ if distortions were impossible, then the scientist would always prefer to choose $a_1 = 0$. We remark that several other papers consider principal agent problems where the agent follows a mixed strategy due to a lack of commitment at the time of contracting; see, for example, Fudenberg and Tirole (1990). Still, to the best of our knowledge, we believe the mechanism isolated here for mixed strategies in this setting is new.

We briefly comment that other forms of scientist private information could be prevelant, and multiplicity may be an issue. For example, if the scientist has private information on θ, the setting becomes a signalling game and the choice of experiment has an additional impact of conveying this private type. The complications arising with these settings, while interesting, are left to future work.

2.5.2 Contractibility

There are many possible interpretations of the benefits to the scientist, and in this paper we prefer an interpretation where these benefits are non-monetary. For example, if drug development is successful, other researchers may view this work as a valuable contribution, and follow-on by working on similar problems. These kinds of payoffs are not easily contractible, and we believe this is a good approximation to many kinds of research activity, particularly research taht is too costly or long term to be undertaken by private entities. On the other hand, for research where the follow-on work is very quick, it may be more feasible to introduce transfer as a function of the chosen experiment. We point out that some kinds of contractibility would make the model degenerate. For instance, if the scientist's experiments could be restricted a priori, then there would be no benefits from changing the transparency regime.[18] Alternatively, if the developer could commit to an effort profile, they could set $e = 0$ unless a given experiment were chosen. Without career concerns, under

[18]This would be difficult if no third party could determine the set of feasible experiments before they were chosen, as in the previous extension.

this policy, the scientist could only possibly obtain non-negative payoff from choosing that experiment. In that case, any experiment that is individually rational for the scientist could be implemented.

However, these comments do not imply that our insights are sensitive to an extreme lack of contractibility. Suppose, for instance, it were possible to pay the scientist an additional amount if they undertake some given experiment, say $\mathcal{I}(\overline{a})$, which is preferred by the developer to the experiment the scientist would undertake with non-transparency. In this case, the scientist would be compelled to undertake $\mathcal{I}(\overline{a})$ if payment for doing so was greater than the difference in the scientist's payoff. Ultimately, the conclusion of Theorem 7 remains valid, even with transfers (though there are cases under which transfers with transparency outperforms non-transparency), since the cost of compelling the scientist to make this change may be larger than the gain that the developer obtains from the more informative experiment.

2.6 Conclusion

2.6.1 Discussion of Model Assumptions

The abstraction of the model is meant to focus on intuition, rather than explicitly mapping to a given setting. Still, one may wonder whether the assumptions of the model make sense. Perhaps most important is the assumption that the transparency regime is set by an outsider over whom the scientist has no control.[19] It could be that the only way of documenting experimental methods is through a pre-registration database, and that the scientist is unable to credibly do this on their own. Why this does appear to be the case in a variety of disciplines is beyond the scope of this paper. It is not difficult to imagine, though, that without a formal mechanism for doing so, credibly demonstrating that the protocols followed were scrupulous might impose prohibitive costs on the scientist.

Our specification that the scientist's benefit comes from the follow-on work of a developer

[19] We mentioned this possibility in Section 2.3.1.

is simply to provide a concrete story within which our results can be presented. We could have alternatively just been concerned about informativeness, and taken as a reduced form the societal objective. In other settings, it may make more sense to think about follow-on researchers extending or elaborating on the scientist's results (with further citations). In that case, one may be interested in modeling the citation process explicitly or describing the "feedback process" through which the scientist, in turn, learns from follow-on researchers. Our analysis is relevant to these settings, but accomodating them explicitly is left to future work.

We imposed some technical assumptions on the space of information structures (e.g., continuity) to ensure equilibrium existence. We also imposed parameterizations on the set of information structures (and restrictions on costs) in order to allow for limited commitment while still being able to interpret the results. While these assumptions are restrictive, we believe that the model formalizes the basic intuition represented in the introduction. Hence the model seems to be an appropriate formalization of the debate over transparency referenced in the introduction. From that perspective, this paper echoes Wilson (2014), on the importance of careful economic reasoning in the design of research environments.

2.6.2 Final Comments

This paper has shown why a sender's inability to fully commit to an information structure may make a receiver better off. While we illustrated the main forces at work through a simple example, our general model clarified that complementarity (which is natural, at least in some settings) between the different kinds of research actions drove the result. As an application, we have considered whether transparency requirements should be more widespread in academic disciplines. The key insight is that non-transparency on one dimension can induce scientists to exert more effort or incur more costs on another dimension, in a way that ultimately makes those interested in the results better off. In assessing this conclusion, we have primarily been concerned with the interest of those who use the results downstream. This is, of course, a simplified view of welfare, but it is motivated by the idea that society invests in scientific research so that it produces results which will be useful for others downstream.

In the main model we assumed that the only policy lever under the designer's control is whether the developer observes the scientist's choice of protocols. For example, a funding agency may be able to implement pre-registration requirements, but might not have the authority to dictate specific steps that scientists follow. In the extensions we considered what might happen if transfers were also available at their disposal, and demonstrated that the main conclusion would not be changed. Still, the main model highlighted the relevant tradeoffs of transparency. On the other hand, transparency also might have an advantage of allowing for richer punishments, something we do not consider here, though they are natural places for future work.

Though our model is theoretical, it still allows us to comment on the active debate on the costs and benefits of transparency mentioned in the introduction. Our main model has highlighted that this debate should consider the extent to which follow-on research influences which experiments scientist choose to perform. In cases where it is significant, our model shows that scientists have a natural incentive to both add informational content to their experiments and bias their experiments. In this case, non-transparency encourages scientists to choose experiments which are inherently harder to bias. If these experiments are particularly difficult as well, then scientists may not be sufficient motivated to undertake them under transparency, but if they are more informative then presumably society is better off if they do.

Finally, we remark that the intuition for many of our results would apply to the research process more generally, and not just academic publication. While we called the person who acts after the scientist the developer, in general this could be any individual who will use the scientist's research, and whose actions the scientist would be interested in influencing. We have shown that in general, a sender might be incentivized to acquire information but not distort a signal, or conversely. We provided a framework which describes when each will happen, so long as there is some sense in which these two types of actions can be distinguished. The observation that non-transparency can be used as a mechanism for money burning will likely have applications beyond the focus of this paper, as these concerns are similar to those that have been raised in a wide variety of contexts in organizational economics.

Chapter 3

Prototyping Under Competition

In the mid-1990s, the Department of Defense (in conjunction with foreign governments) requested that Boeing and Lockheed Martin each begin the development a fighter jet that would combine the functionality of several aircraft being used by different branches of the U.S. military, as well as the U.K. Royal Navy and other American allies. The program, referred to as the Joint Strike Fighter (henceforth JSF) due to its intended use by different branches working in tandem with one another, was one of the most (if not the most) ambitious technological procurement programs in recent military history. Lockheed Martin was awarded the right to develop the aircraft in part due to the ability to perform Short Takeoff and Vertical Landing (STOVL). However, several years after the contract was awarded, it became clear that the aircraft was 2 % above the target weight, throwing the entire program into jeopardy.[1] The estimated cost of the project ended up exceeding a *trillion* dollars, and while opinions of the program vary substantially, the dominant view (at least among policymakers) is that the end result was a blunder for both the government (due to the enormous cost, for a project that was in part justified by potential for savings), as well as Lockheed Martin (due to the pressure they faced to make the program work).

This paper studies the problem of allocating contracts for the purposes of developing

[1]See Blokcom (2003) for a review of the program and issues involved. For a more recent assessment of the difficulties facing the JSF program and policy recommendations prescriptions meant to remedy them, see GAO (2016).

and subsequently procuring new technologies, as was done for the JSF. At an abstract level, there are many possible ways to allocate such a contract; for instance, the principal could ask each participant for detailed descriptions of every feasible project. In practice, the government tends to award contracts based on a contractor's demonstrations or descriptions of it.

I refer to this method of allocating contracts as *prototyping*. Prototyping is by no means limited to the military, and even within the military there is enormous variation in how prototyping is done, in part due to the variation in the size of contract awards. Indeed, the JSF was extraordinary in its size and profile, and therefore arguably not representative of the vast majority of procurement activity (despite its clear significance). But while there are a relatively small number of major weapons systems, there are a very large number of smaller contracts which are awarded regularly for the purpose of technological development.[2] In this paper, I take a broad view of what prototyping entails, but am primarily interested in the following kinds of settings:

- The contract award has value for the procurer as well as the developer (due to the ability to develop more in the future, the acquisition of knowledge, etc.).
- Residual uncertainty remains even after the contract has been awarded, and where eventual infeasibility of the prototype is costly for all parties.

These apply to a large number of acquisition programs and are highlighted by the above discussion of the JSF. Within these settings, I am interested in allocation mechanisms consisting of an agreement to develop a particular project (i.e. the prototype), based on a description of it. Allocating in this way is restrictive, and rules out several allocation mechanisms which many theoretical models may implicitly assume are feasible, but which seem less common in practice. For instance:

- Awarding the contract to a participant, and subsequently dictating which particular

[2] For instance, see Bhattacharya (2016) for an empirical analysis of a particular method the Department of Defense uses to allocate contracts to small businesses.

project they undertake (unrestricted by the initial prototype)

- Asking for participants to submit a set of prototypes, and auctioning off the contract based on the willingness to do each one

This paper asks why a principal may not opt for these more sophisticated mechanism. I develop a deliberately simple model in order to study prototyping formally. In my model, agents are endowed with prototypes, observing privately some characteristics and subsequently making a report to the principal. The principal then allocates the contract to one of the agents, and payoffs depend on the prototype chosen and a subsequent event that depends on the agent's private information. I allow for rich communication protocols, but am interested in conditions under which optimal communication mechanism will simply entail each agent suggesting which prototype they would like to develop, and doing so whenever they are chosen. I refer to these as *prototyping mechanisms*.

I am interested in the interaction between competition and the use of prototyping mechanisms. In practice, determining how much to solicit competition is often the *main* relevant policy lever. For example, Stage I projects in the SBIR program are often restricted in terms of contract duration and monetary award, but there is substantial variation in how many participants are selected to participate (see Bhattacharya (2016)). Indeed, the Competition in Procurement Act of 1984 significantly restricted the ability of government organizations to allocate contracts without "full and open competition" (see Manuel (2010) for more details on competition in procurement for federal agencies). Certainly there are compelling reasons to be interested in maintaining competitive procurement, many of which are not modelled here. Instead, this paper focuses on one particular reason that competition is viewed as per se beneficial: That if enough prototypes are present, the probability of finding a particularly good one is high. This paper shows that this logic is incomplete, and in fact it can backfire if not done effectively.

Specifically, my results show that when incentives are aligned, the principal can implement the optimum by simply asking each agent to suggest a prototype and subsequently asking some agent to develop it. In other words, prototyping mechanisms achieve the optimum. A contribution of this paper is to distinguish between deterministic project proposal and deterministic agent selection. While prototyping mechanisms are defined by their use of deterministic project propsal, they are *not* necessarily deterministic since the selected agent can be chosen stocahstically. The alignment assumption introduced may seem strong, since it forces a deterministic optimum in the problem with only a single agent. However, I show by example that a deterministic optimum in the single agent benchmark does not by itself imply that prototyping mechanisms are optimal in competitive settings. In other words, without any restrictions on preferences, prototyping could be optimal in non-competitive settings but suboptimal in (otherwise similar) competitive settings.

I proceed to provide several results describing the nature of competition under prototyping mechanisms. First, if the aforementioned randomizations can be committed to, then I demonstrate that competition is beneficial if it is costless to obtain additional participants. In contrast, if randomizations are not possible, then it is possible for competition to be counterproductive. I also demonstrate that the results are maintained when effort must be undertaken, though this requires relaxing symmetry assumptions. Finally, I consider the setting with fines, showing that randomization may still be necessary if the fines are not so large as to break the alignment between principal and agent.

I continue with the model in Section 3.1, and provide a simple parameterization to highlight the key insights in Section 3.2. I then walk through the first best benchmark in Section 3.3, with the main results on the optimality of prototyping mechanisms being in Section 3.4. I highlight the role of commitment in Section 3.5, with some other extensions being discussed in Section 3.6. Literature is left to Section 3.7, concluding in Section 3.8. Proofs and additional results are in the Appendix.

3.1 Model

A principal wishes to hire one of N agents to develop a new technology. Agent $i \in \{1, \ldots, N\}$ is endowed with a set of possible prototypes. A prototype is denoted x_j, with $x_j \in \{x_0, \ldots, x_{M_i}\}$, for $M_i \geq 1$. I refer to x_0 as the *safe* prototype, which I assume can be made by all participants and is identical. All other prototypes are referred to as *risky*.

For each risky prototype x_j that an agent is capable of producing, the agent observes some $\theta^i_j \sim G^i_{x_j}$. The vector $(\theta^i_j)_{1 \le j \le M_i}$ is privately observed by each agent. However, the set of possible prototypes $\{x_0, \ldots, x_{M_i}\}$ for each agent i is commonly observed by all parties, as are the distributions $G^i_{x_j}$. Additionally, I assume that $G^i_{x_j}$ is finite.

The principal allocates the contract according to a mechanism described in the following subsection. After the allocation is made, an outcome occurs, which I denote by $y \in \{0, 1\}$. The event $y = 1$ is referred to as *success* and the event $y = 0$ is referred to as *failure*. The safe prototype results in success with probability 1, but risky prototype x_j results in success with probability θ^i_j.

3.1.1 Mechanisms

I introduce two classes of mechanisms that the principal may be able to use. First, a *general mechanism* allows each agent i to make a report $m_i \in M_i$. The mechanism then chooses an agent, as well as a project for them to develop, according to some rule $X : M_1 \times \cdots \times M_N \to \Delta(\cup_i \{x_0, \ldots, x_{M_i}\})$ (i.e. it chooses an agent and a prototype for the agent).

I do not allow for transfers to be made between any parties in the main model.[3] By the revelation principle, it is without loss of generality to assume that each M_i is each agent's type space, and an equilibrium is played where all agents report their types truthfully. However, this results in a difficult set of mechanisms to work with.

I refer to a simpler set of mechanisms as *prototyping mechanisms*. These are mechanisms where first, the principal determines a *set* $D_i \subset \{x_0, \ldots, x_{M_i}\}$, and each agent chooses some $x \in D_i$. The principal then chooses an agent according to some rule $Y : \prod_{i=1}^N D_i \to \Delta(\{1, \ldots, N\})$. As mentioned in the introduction, allocation of contracts for new technologies in practice involve simple descriptions of the technology to develop, and do not necessarily involve detailed reports of all alternative technologies (and in fact, details of the given

[3]This is common in early stage prototyping. In practice, research budgets are small, and most of the incentives are derived from the ability to win a contract in the future.

technology are often difficult to verify). Hence while general mechanisms are the more theoretically appealing concept, prototyping mechanisms appear to be more common in practice. However, I am not directly able to appeal to the revelation principal for prototyping mechanisms.

3.1.2 Payoffs

Any agent that is not selected obtains a payoff of 0. If the outcome is y, then an agent obtains payoff $v_A^y(x)$ and the principal obtains payoff of $v_P^y(x)$. I assume throughout that all parties have higher payoffs following a success than following a failure: $v_i^1(x) \geq v_i^0(x)$, for all x and $i = A, P$.

As the name suggests, I interpret $y = 1$ to denote the case where the prototype is successfully developed and the event $y = 0$ to denote the case where it is not. In practice, parties may in fact have positive payoffs in the event of failure. For instance, knowing the difficulties of some projects makes later projects more efficient, or may still result in improvements on a known technology.

At times, I will be interested in cases where the following assumption holds:

Assumption 1 (Alignment)**.** *Whenever* $x_j \in argmax_j\{\theta_j^i v_A^1(x_j) + (1-\theta_j^i)v_A^0(x_j)\}$, *it also holds that* $x_j \in argmax_j\{\theta_j^i v_P^1(x_j) + (1-\theta_j^i)v_P^0(x_j)\}$.

Alignment forces both principal and agent to have the same maximizing prototype choice conditional on private information. When this assumption holds, I will refer to the prototype that is optimal for both principal and agent as the *efficient prototype*.

While alignment may seem to be a strong assumption, it will hold whenever surplus between the principal and the agent is determined by an ex-post Nash bargaining protocol over surplus. A similar kind of surplus division is assumed in Bhattacharya (2016), and also appears reasonable based on conversations with practitioners. In the Appendix, I consider a version of the model where the agent observes a cost draw after observing whether the prototype is feasible, and where prototypes are distinguished by their value to the principal. In this setting, whenever the cost has a linear virtual value, the alignment assumption will hold (using the Myerson optimal transfers given the principal's value) as well. Without linearity, it can also hold for discrete type spaces, and I also show how the concavity of the virtual value influences the incentives to overstate or understate a project's quality.

3.1.3 Summary and Discussion

To summarize, the model consists of:

- Agents observing feasible prototypes and private information about their quality.
- The principal designing some announcement mechanism.
- Agents choosing their announcements, and the contract being allocated according to the principal's rule.
- The selected agent develops their prototype, the ultimate feasibility being realized after the agent is selected.

I am assuming that once the agent is selected, it is impossible to "go back" and select another agent. For instance, it was practically impossible for the JSF to be reallocated, and there is often significant red tape when there are attempts to bring back "failed" projects. On the other hand, both parties may prefer a failed project to no project at all. The important driver of incentives in the model is that agents have an additional incentive to be chosen, whereas the principal may be indifferent between various agents conditional on their private information.

Another key assumption relates to the independence of private information across protoypes (and agents). This assumption is appropriate in order to best scrutinize the common reasoning for the use of competition, as a means to getting more "good draws." If quality were correlated across agents, then this impact would be diminished. When dynamics are significant, however, this would be a more problematic assumption, as looking at past performances may dictate what kinds of inference about feasibility that the principal would make in the future.

3.2 An Example

The following example illustrates many of the key ideas in the more general analysis (although takes advantage of some results that are still to come). It demonstrates that the optimal prototyping mechanism involves randomization and highlights the reasons this randomization arises. Specifically, we consider the case where there are two prototypes (one safe and one risky) and where the risky prototype succeds for agent i with probability $\theta_i \sim F$. The risky prototype, if successful, gives a payoff of 2 to both the winning agent and the principal if selected, but gives a payoff of 0 to both if it is not successful. The safe project, by contrast, gives a payoff of 1 to both for sure.

Note that the alignment assumption holds in this setting. With only one agent, since the principal and the agent both have the same utility function, it is clear that the agent implements the principal's favorite prototype even in the second best. Specifically, the principal's optimal policy consists of a *threshold rule* whereby the safe project is implemented if $\theta_i \leq 1/2$, and the risky project is implemented otherwise.

In contrast, a deterministic threshold rule may not be optimal in the case with two agents. Note first that if each θ_i were observed by the principal, then having more agents participate is clearly beneficial—the principal faces two draws instead of one, and hence the best θ_i is more likely to be higher. To make our point simplest, suppose that F is supported on two values, $0 < \theta_L < \frac{1}{2} < \theta_H < 1$, with $\mathbb{P}[\theta_H] = q$. The first best in this case would (1) have the agent develop their preferred prototype if they are selected, and (2) choose the agent with the higher θ_i if the θ_i differ. Consider a direct revelation mechanism where the principal uses this mechanism, and let us conjecture that agents follow truth-telling. In this case, an agent of a low type knows that with probability q he will not get the award, since this is the probability the other agent has a higher type. Let $y_{i,t}$ be the probability that agent i is allocated the contract if both are of type t. In this case, incentive compatibility requires:

$$(1-q)y_{i,L} \geq (1-q+qy_{i,H})2\theta_L$$

Adding this up across both agents and noting that the first best always allocates the object:

$$(1-q) \geq (2-q)2\theta_L \Rightarrow \frac{1-q}{2(2-q)} \geq \theta_L.$$

When $q = 0$, the principal can implement the first best by simply allocating the safe prototype to one of the agents. But for $q \gtrapprox 0$, the first best requires allocating the risky project with positive probability, which may not be possible if the risky prototype is always developed over the safe prototype; the condition for this to occur becomes $\frac{1}{4} \gtrapprox \theta_L$. So whenever $\theta_L > 1/4$, even if it is very unlikely that the other agent is a higher type, agents would still choose the (less preferred) risky prototype. Hence overpromising can still be a concern, even if it is extremely unlikely that the competitor is actually able to implement a better prototype.

Now, the left hand side of the inequality is decreasing in q, so in fact when q is higher truth telling is more demanding. So, overpromising is indeed less of a concern when the probability of a successful risky prototype is smaller. However, when q is smalller, the gains to deviating from truth telling are higher as well. That is, if it is very likely that the other agent is a low type, then it is also very likely that lying results a win. While it results in a suboptimal choice, it is still preferred to losing, which is a real risk from telling the truth. Hence the incentive to overpromise does not vanish in the limit.

This outcome highlights a potential for a *perverse effect of competition*. That is, suppose $\theta_L > \frac{1-q}{2(2-q)}$, and that the equilibrium played is "proposing the risky prototype." [4] In that case, the payoff to the principal is $2(q\theta_L + (1-q)\theta_H)$. If $q\theta_L + (1-q)\theta_H < 1/2$, then the result is actually worse than if there were only one agent.

Note that competition can still be beneficial, even if agents develop their preferred prototype, by utilizing a stochastic mechanism that randomizes across agents. Indeed, suppose the mechanism is symmetric, and let $y_{i,L} = y_{i,H} = 1/2$. Let z denote the probability that the winning prototype is safe if both prototypes are proposed. One need only choose z

[4] This is in fact the unique equilibrium.

so that incentive compatibility is satisfied:

$$\frac{1-q}{2} + qz \geq 2\theta_L \left((1-q)(1-z) + \frac{q}{2} \right)$$

which occurs if:

$$z \geq \frac{2\theta_L(1-q) + q\theta_L - \frac{1-q}{2}}{q + 2\theta_L(1-q)}$$

As $q > q\theta_L - \frac{1-q}{2}$, this condition can hold with equality while maintaining $z < 1$.[5] Setting z to be equal to the right hand side delivers the optimal mechanism, which indeed does better with two agents than one agent. Notice, however, that any agent that proposes the safe project delivers the principal a lower utility. Hence setting $z > 0$ requires commitment on the part of the principal to allocate a contract to a project that delivers a lower utility in expectation. If this commitment power were absent, then indeed competition could be harmful.

3.3 Hard information benchmark

In this section, I consider the case where each agent's private information is verifiable, so that the allocation rule can depend on the agent's realized type vector. This corresponds to the case where all information regarding the prototype is observable and contractable. In that case, the principal can simply ask all agents to report their entire vector of private information, and hence the principal is able to pick the prototype that maximizes their payoff. While this competition results in a better prototype being picked with higher probability, it also damages the payoffs that each agent obtains. This discussion yields the following proposition, whose proof is omitted.

Proposition 0: *With observable private information, the principal's payoff from an optimal mechanism is increasing in the number of agents, while the average payoff for each agent is decreasing.*

[5]Note that as $q \to 0$, $z \to 1 - \frac{1}{4\theta_L}$.

Even with observable information, the principal may still (for whatever reason) prefer mechanisms that are simple. For example, it may be difficult to get agents to willingly announce information even if it is hard. In this setting, without the alignment assumption, the principal may do better by conditioning the allocation on information the agent would not voluntarily disclose; for instance, if a certain prototype was always less preferred to the safe prototype by the agent, but preferred by the principal for some agent types. However, when the alignment assumption holds, these incentives disappear, with or without competition. So in terms of implementation, when the alignment assumption holds, the principal can restrict to (fully deterministic) prototyping mechanisms. I record this observation as follows:

Proposition 0′: *Under the alignment assumption with verifiable agent types, the optimum is implemented by asking each agent only for their favorite prototype and the corresponding quality parameter.*

Under the alignment assumption, asking each agent for their favorite prototype implements the optimal mechanism, no matter how many agents there are, and hence the project selection mechanism does not depend on the number of agents. Without the alignment assumption, the principal's outside option when not selecting a given agent is increasing, meaning that agents are more likely to propose projects that are favorable to the principal. And since information is non-falsifiable, this shift in incentives cannot hurt the principal. However, even without the alignment assumption, there is no incentive for the principal to randomize over projects, provided the mechanism can elicit enough information: given any report, the principal can pick the prototype that is optimal for them.

3.4 Prototyping

This section provides conditions for the optimality of prototyping mechanisms. I emphasize the role of the alignment assumption, and show that some version of it is necessary to explain the widespread use of prototyping. First, I show that in general, one can look at a restricted set of general mechanisms without hurting payoffs:

Lemma 6. *The optimal general mechanism can be achieved via a policy under which the project an agent develops depends only on the agent's own reports.*

The proof of this result is straightforward and relegated to the appendix. The idea is to replace any rule where the randomization of each other player's reports are "simulated out." The proof demonstrates that, due to the structure of payoffs, such a replacement is payoff equivalent for all parties, even if it may lead to a different outcome.

Given this lemma, it is worth comparing the case of a single agent to the case of multiple agents, where the principal need not worry about which agent to select. Difficulties in the single agent problem are well known due to the difficulty in applying the revelation principle with deterministic mechanisms, and it is also known that certain generalizations to multi-agent contexts are not feasible, as highlighted by Strausz (2003). The lemma points out that under certain separability assumptions, randomizations can be done across agents. One may therefore wonder whether a deterministic single-agent problem at least guarantees a deterministic prototype choice in the multi-agent problem.

Consider the following example:

Example 1. I modify the example from Section 3.2 slightly. As before, suppose the safe prototype delivers payoff of 1 for all parties and the risky prototype delivers payoff of 2 for all parties, with the same set of possible types for the agents. However, now also assume that there is a third prototype that can be developed, which yields a payoff of $1 - \varepsilon$ to the principal, and $-D(\varepsilon)$ to the agent.

Clearly the single agent problem has a deterministic solution, and coincides with the case in Section 3.2. There, I also showed that the optimal mechanisms without the third prototype present involves randomization across agents, but not across projects. With the third prototype, it is never optimal for the third prototype to be chosen if the principal restricts to prototyping mechanisms (i.e., deterministic project choice): both principal and agents do better by choosing the safe prototype over the third prototype. The optimal mechanism from the example is therefore again optimal when there is deterministic prototype choice (even when randomizations are done across agents).

Consider instead the following mechanism: If both agents announce $\underline{\theta}$ then choose the safe prototype, randomly allocated between the two agents, and if both agents announce $\overline{\theta}$ then choose the risky prototype, randomly choosing between the two. However, if one agent announces the safe prototype and one announces the risky prototype, allocate the contract to the agent that announces the risky prototype with probability 1, but have them develop the third prototype with probability r.

In this case, an agent of type θ_L prefers to announce the safe prototype if:

$$\frac{1-q}{2} \geq 2\theta_L((1-q)(1-r)+\frac{q}{2})-(1-q)rD(\varepsilon),$$

and is indifferent if:

$$r = \frac{3\theta_L - 1/2}{2(D(\varepsilon)+2\theta_L)}.$$

Whenever this indifference holds, a type θ_H will strictly prefer to announce the risky prototype. Taking $D(\varepsilon) \to \infty$ shows that it is possible to also have $r \to 0$. However, taking $\varepsilon \to 0$ simultaneously demonstrates that the principal obtains first best payoffs in the limit, and hence does strictly better for some $D(\varepsilon), \varepsilon$. ■

The example shows that prototyping mechanisms may be optimal in single-agent problem, but not in the corresponding multi-agent problem where agent types are IID. The reason is simple; it is less costly for the principal to randomize over projects than randomize over agents. Therefore, a stronger assumption is needed in order to guarantee the optimality of deterministic project choice (and hence prototyping mechanisms). This is the role of Assumption 1. The point that randomization can be part of the optimal contract in order to act as a form of money burning is made by Kovac and Mylovanov (2009). The example shows that the presence of competition gives one channel for this to arise.

However, under the alignment assumption, I am able to show that with a single risky prototype, it is without loss to consider mechanisms which restrict to prototype announcements:

Lemma 7. *Under Assumption 1, the optimal general mechanism involves a deterministic project* *choice, and hence can be implemented as a prototyping mechanism, whenever each $G^i_{x_j}$ has support of cardinality 2.*

The idea is that replacing any arbitrary mechanism with one that implements the mutually preferred allocation strictly increases the surplus of all parties. Without the alignment assumption the result will generally fail, as the above example shows. Intuitively, alignment ensures that the only way of harming the agent is by allocating the contract to the other agent. Note that this is not present in Example 5.1, since there is disagreement between the principal and the agents over the third prototype.

Next, I demonstrate the natural intuition that it is the "bad types" that wish to mimic the "good types."

Lemma 8. *Whenever the incentive compatibility constraints are binding, the safe type IC constraints are binding in any optimal mechanism.*

This lemma can be used to calculate the optimal mechanism in examples, since it is pinned down by the safe type not wanting to mimic the risky type.

Given that the optimal mechanism does not always allocate the contract to the ex-post efficient prototype, one may wonder whether competition actually has any benefits. It turns out the answer is in the affirmative:

Proposition 9. *Suppose all agents are symmetric, with one risky prototype per agent. Then there exists an optimal symmetric mechanism, and the principal's principal's payoff is increasing in the number of agents, but is bounded as the number of agents grows large.*

The fact that the principal's payoff is bounded follows immediately from the observation that even if a risky prototype could be guaranteed, the principal's payoff would be bounded (hence the proof is omitted from the appendix). It follows that while the profit is increasing in the number of agents, there are diminishing marginal returns to additional agents.

It is worth comparing this result to Mylovanov and Zapechelnyuk (2017). They consider a setting where a princpial can impose (exogenous type dependent) fines on agents if they misreports, though imposing these fines is costly for both the principal and the agent. Their result is that competition has limited benefit, in that there is a finite number of agents above which the principal's payoff does not increase.

The intuition behind the basic structure of the optimal contracts in their paper is similar to ours: The principal is limited in how much they can deter misreports, and hence does not always allocate the contract to the agents that have the best type. However, in their model, the cost of punishment is exogenous. Here, it is endogenous due to the fact that another agent can be awarded the contract, and one that potentially has a better prototype that they can develop.

3.5 Commitment Assumptions

The previous analysis showed that under restrictions on the environment, an optimal mechanism generally must allow for randomizations. As pointed about by Laffont and Martimort (2002), this necessitates a certain degree of commitment to an outcome that is ex-post suboptimal, and hence requires credibility for the principal to carry out this randomizations. Circumventing for this kind of commitment problem may explain why, for example, an organization may give the power to allocate rights to a separate division with clear guidelines (as is done with SBIR as pointed out by Bhattacharya (2016)). However, it may also be unavoidable for cases like the JSF, which had high stakes and flexible criteria, so that any explicit randomization would have likely been infeasible.

In this section, I consider a benchmark where the principal is *restricted* to utilizing prototyping mechanisms, and *cannot* commit to a particular randomization. That is, I impose the following timing:

- Agents all submit prototype recommendations.
- The principal then chooses a prototype to develop, among the list of submitted prototypes.
- The prototype is developed, with the outcome y being realized and payoffs being achieved.

Proposition 10. *Let $N = 2$ and suppose the principal cannot commit to randomizing after prototypes. Let the alignment assumption hold with a single risky prototype for each agent. Suppose further that in the optimal contract with commitment, the contract is allocated to the safe prototype over a risky prototype with positive probability. Then the optimal mechanism involves the principal preventing one of the agent's from participating.*

Note that with 2 agents, there may exist an equilibrium which delivers the same payoff as the single agent mechanism: Simply choose one of the agents and treat the other agent's report as babbling. In order for this to be an equilibrium, however, it must be that without any information, the safe prototype is preferred to the risky prototype. Considering the example from Section 3.2, if q is very high, then it cannot be an equilibrium to treat one agent's reports as truthful and the other as babbling, since otherwise the principal would profitably deviate by choosing the babbling agent. Hence in this case, competition may indeed strictly hurt. Furthermore, there is always an equilibrium under which the strategy is pure babbling. Hence under the alignment assumption, the payoffs to the principal would alsways strictly decrease with the addition of a second agent if such an equilibrium were played.

3.6 Extensions

3.6.1 Fines

The main model assumes that there are no transfers available. For early stage progress, this is often the case; for instance, the SBIR program provides firms with small lump sum loans if they are awarded a first stage contract, with the primary incentives often being the possibility of contract awards. However, these awards often originate from a completely separate entity (e.g. the Navy) and hence are not explicitly part of the original contract (see Bhattacharya (2016) for more institutional details).

Still, one may consider cases where the outcome y is at least partially contractable, and where liability constraints are not binding so that some payoffs can be contracted on ex-ante.

With large scale projects, for instance, the firms often have multiple projects, and hence the ability to transfer payoffs between principal and agent should arise in some form. [6] Can fines enable competition to be beneficial when the principal is not able to commit to randomization?

It is not hard to see that the answer is sometimes yes, but it turns out the ability to fine the agent has limited ability to restore the benefits of competition when commitment is not present. I demonstrate this observation by revisiting the simple example from Section 3.2. Suppose the principal were able to enforce a fine against the agent in the case of failure instead of distorting the allocation. To induce truth telling, in a symmetric mechanism, a fine would need to satisfy:

$$(1-q)(1/2) \geq (1-q+q(1/2))(2\theta_L - f(1-\theta_L)) \Rightarrow f \geq \frac{\theta_L(4-2q)-(1-q)}{(2-q)(1-\theta_L)}.$$

One should read this expression cautiously however. For example, suppose that $q = 1/2, \theta_L = 2/5$ and $\theta_H = 3/5$. Based on the above analysis, the overpromising effect arises, and a fine that rectifies it is $7/9$. Suppose such a fine could in fact be imposed on the agent, and that the principal collects the payment. In this case, one can check that the principal would actually be better off[7] limiting discretion and forcing both agents to pick the risky prototype! Without the agents making any choices and the allocation being random, clearly the principal's payoff with one agent coincides with the solution for two agents. So while indeed fines can be imposed to mitigate overpromising, even in a world with fines, competition might only have limited benefits.

[6]In conversations with practitioners, the dominant case is that since the major defense contractors have mulitple projects over a long horizon, a perceived misrepresentation on one project would be corrected for on another project. Hence payoffs would adjust, even though such adjustments would not be made explicitly as part of the contract. Such a story is reminiscent of the relational contracts literature; making this connection explicitly is intriguing, but beyond the scope of the present paper. However, see Board (2011) for a model along these lines.

[7]Albeit only very slightly.

3.6.2 Moral Hazard

In practice, prototyping involves effort on behalf of participants, and the incentives to create a future project with high surplus is a significant driver of incentives. Indeed, much of the existing literature (Che and Gale (2003), Taylor (1995), Fullerton et. al. (2001)) focus on this aspect of R & D. I comment on how this possibility influences the optimal mechanism.

Assume that there are N agents and each has access to only a single risky prototype. Prior to entering the mechanism, I assume that each player chooses an effort level $e \in E$ at cost $c(e)$, normalizing the cost of the lowest effort level to be 0. The parameter $\theta \sim G(e)$ is drawn, as a function of the effort chosen. For simplicity, I use the setting of the example, where $\theta \in \{\theta_L, \theta_H\}$ with $\mathbb{P}[\theta_H] = e$.

The presence of moral hazard adds an additional cost of randomization. In the single agent case, the agent chooses e to maximize:

$$-c(e) + 2\theta_H e + (1 - e)$$

In contrast, in the multiple agent case, with the optimally chosen symmetric allocation probability, e_i maximizes:

$$-c(e_i) + e_i \cdot (\theta_H(e_{-i} + (1 - z)(1 - e_{-i})))) + (1 - e_i)\left(\frac{1 - e_{-i}}{2} + e_{-i}z\right)$$

As is typical, agent efforts are strategic substitutes. Hence the presence of the second agent lowers the incentive for the first agent to exert effort.

If E is discrete, then Proposition 9 can be used to show that adding additional identical participants can the principal whenever all agents have strict incentives to exert their given effort level. If the agents are in fact indifferent between a given effort level and a lower one, then adding another agent changes the distribution, which may hurt the principal. On the other hand, assuming the principal can choose arbitrary mechanisms, adding additional agents cannot hurt, and adding sufficiently many agents will help in the case that exerting minimal effort can still result in θ_H being drawn.

However, unlike in the case without moral hazard, the environment cannot necessarily be symmetrized to ensure that the moral hazard constraints are still satisfied. Indeed, symmetry assumptions may be unnatural simply because there may not be an equilibrium where both agents exert the same effort level (due to strategic substitutability). If the principal is restricted to choosing symmetric mechanisms, therefore, competition may indeed lower the principal's payoffs.

3.7 Literature Review

Our main analysis studies delegated project choice. A single agent version of this problem is studied by Armstrong and Vickers (2010). Bar and Gordon (2014) consider a general analysis of project selection with multiple agents, but allow for transfers under limited liability. An application of delegated project choice to merger review (with multiple agents) is the focus of Nocke and Whinston (2014), who look at restricted mechanisms motivated by similar issues with eliciting information by anti-trust authorities. Rantakari (2014) also studies multi agent project selection, but also assumes agents obtain benefits from the other parties, and does not focus on the role of randomization.

Outside of the project choice literature, multiple agent delegation models have been studied by Ambrus, Baranovskyi and Kolb (2017), as well as Martimort and Semenov (2008) for an application to political lobbying. Ambrus, Baranovsky and Kolb (2017) assume that the state is common across all parties (whereas agents get a noisy signal of it) in contrast to the model here, and also only consider the case of commitment under a restricted set of mechanisms. Martimort and Semenov (2008) also consider somewhat different functional forms, although they allow for independent types across agents. However, they restrict to studying deterministic mechanisms, leaving the question of the benefits to randomization open (but also commenting as I do above that this may be difficult to enforce). A related literature that has focused on cases of multiple agents relates to information production by experts in organization. For instance, Dewatripont and Tirole (1999) study advocacy in organizations, and also consider the importance of falsifiability, but model it in a different way. Since agents in their model do not have intrinsic preferences over prototypes, the case of completely soft information is degenerate.

Mechanism design with verifiable information originates with the seminal work of Green and Laffont (1986), who illustrate difficulties with the revelation principle in settings with partially verifiable information. As stated above, Mylovanov and Zapechelnyuk (2017) provide a similar model where the principal can punish the agents ex post for lying, and show that this provides a rationale for random allocation. However, they do not study the delegation problem per se, and restrict preferences more substantially than I do here (and hence do not comment on how these assumptions influence the incentives to randomize). Additionally, as stated, I arrive at a different conclusion regarding competition (since there is always some positive gain from additional participants). Ben-Porath, Dekel and Lipman (2014) allow for the principal to potentially check the type of the agent, which entails an ex-ante (instead of ex-post) verification. Halac and Yared (2017) study a delegation setting where the principal can verify an agent's private information, and Silva (2017) studies this question in a criminal justice setting.

3.8 Conclusion

This paper considers optimal project selection mechanisms, and provides conditions under which they can take the form of prototyping mechanisms that are often seen in practice. I demonstrate how competition in delegation settings can lead to the necessity of randomization, and that under some assumptions on the underlying environment these randomizations can be taken to be over the agent selected (i.e. allowing for deterministic project choice). I also highlight in a simple model how the benefits of competition are connected to commitment, thereby demonstrating that more need not help if this is absent.

Many aspects of reality are missing from the main model. For instance, I take a stark perspective on the verifiable/non-verifiable aspects of projects, but in practice there may be other ways of partially verifying an agent's private information. For instance, the principal may be able to conduct audits (as in Halac and Yared (2017)), or there may be costs associated with falsifying the private information. Therefore, it remains open how much unverifiability is necessary for the results. Additionally, one may endogenize the outcome after failure by allowing for explicit dynamics, for example. Dynamics are particularly interesting since agents may be learning about the feasibility of their projects, or may have the option to sit out of a given round and delay. De Clippel, Eliaz and Rozen (2017) provide a model that studies some of these issues, where a principal chooses among agents with noisy ex-post information of their type, and where a decision is to be made in every period. Indeed, as stated above, the future contracts are a major driver of incentives in these applications. So while this may be complicated, future work should take these incentives into account explicitly, instead of treated in the reduced-form way that is done so here.

References

Philippe Aghion, Mathias Dewatripont, and Jeremy Stein. Academic freedom, private-sector focus, and the process of innovation. *RAND journal of Economics*, 39(3):617–635, Autumn 2008.

Attila Ambrus, Volodymyr Baranovskyi, and Aaron Kolb. A delegation-based theory of expertise. Technical report, Duke University, May 2017.

Attila Ambrus and Georgy Egorov. Delegation and nonmonetary incentives. Working paper, Northwestern University, Kellogg, November 2014.

Isaiah Andrews and Max Kasy. Identification of and correction for publication bias. Technical report, Harvard University and Massachusetts Institute of Technology, March 2017.

Charles Angelucci. Motivating information acquisition through risky decision-making. Working paper, Columbia Business School, July 2014.

Rossella Argenziano, Sergei Severinov, and Francesco Squintani. Strategic information acquisition and transmission. Working paper, June 2014.

Mark Armstrong and John Vickers. A model of delegated project choice. *Econometrica*, 78(1):213–244, January 2010.

Pierre Azoulay, Alessandro Bonatti, and Joshua Krieger. The career effects of scandal: Evidence from scientific retractions. Discussion paper, MIT, May 2015.

Kyle Bagwell. *The Economic Analysis of Advertising*, volume 3. Elsevier, North Holland, 2007.

Talia Bar and Sidartha Gordon. Optimal project selection mechanisms. *American Economic Journal: Microeconomics*, 6(3):227–255, 2014.

Elchanan Ben-Porath, Eddie Dekel, and Barton Lipman. Optimal allocation with costly verification. *American Economic Review*, 104(12):3779–3813, December 2014.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, March 2015.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. Informationally robust optimal auction design. Technical report, Working Paper, Yale University, University of Chicago and Princeton University, December 2016.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, January 2017.

Dirk Bergemann and Ulrich Hege. The financing of innovation: Learning and stopping. *RAND journal of Economics*, 36(4):719–752, Winter 2005.

Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146:2527–2543, 2011.

Vivek Bhattacharya. An empirical model of r & d procurement contests: An analysis of the dod sbir program. Technical report, Massachusetts Institute of Technology, 2017.

Andrew Biehl. Durable-goods monopoly with stochastic values. *Rand Journal of Economics*, 32(3):565–577, Autumn 2001.

David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272, June 1953.

Christopher Blokcom. Joint strike fighter (jsf) program: Background, status, and issues. Technical report, Congressional Research Service, The Library of Congress, 2003.

Simon Board. Durable-goods monoply with varying demand. *Review of Economic Studies*, 75(2):391–413, April 2008.

Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Dynamic monopoly pricing and herding. *RAND Journal of Economics*, 37(4):910–928, Winter 2006.

Subir Bose, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. Monopoly pricing in the binary herding model. *Economic Theory*, 37: 203–241, 2008.

Jeremy Bulow. Durable-goods monopolists. *Journal of Political Economy*, 90(2):314–332, April 1982.

Rene Caldentey, Ying Liu, and Ilan Lobel. Intertemporal pricing under minimax regret. *Operations Research*, November 2016.

Vinicius Carrasco, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira. Optimal selling mechanisms under moment conditions. Technical report, Working Paper, PUC-Rio, April 2017.

Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, February 2015.

Gabriel Carroll. Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488, March 2017.

Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, September 2013.

Yeon-Koo Che and Ian Gale. Optimal design of research contests. *American Economic Review*, 93(3):646–671, June 2003.

Yiwei Chen and Vivek Farias. Robust dynamic pricing with strategic customers. Technical report, Working Paper, Renmin University and MIT, December 2016.

Kim-Sau Chung and Jeffrey Ely. Foundations of dominant-strategy mechanisms. *Review of Economic Studies*, 74:447–476, 2007.

Lucas C. Coffman and Muriel Niederle. Pre-analysis plans have limited upside, especially where replications are feasible. *Journal of Economic Perspectives*, 29(3):61–80, Summer 2015.

John Conlisk, Eitan Gerstner, and Joel Sobel. Cyclic pricing by a durable goods monopolist. *Quarterly Journal of Economics*, 99(3):489–505, August 1984.

Vincent P. Crawford and Joel Sobel. Strategic Information Transmission. *Econometrica*, 50(6):1431–1451, November 1982.

Jacques Cremer. Arm's length relationships. *The Quarterly Journal of Economics*, 110(2):275–295, May 1995.

Rahul Deb. Intertemporal price discrimination with stochastic values. Technical report, Working Paper, University of Toronto, February 2014.

Geoffroy de Clippel, Kfir Eliaz, and Kareen Rozen. The silent treatment. Technical report, Brown University and Tel-Aviv University, August 2017.

Mathias Dewatripont, Ian Jewitt, and Jean Tirole. The economics of career concerns, part i: Comparing information structures. *Review of Economic Studies*, 66(1):183–198, January 1999.

Mathias Dewatripont and Jean Tirole. Advocates. *Journal of Political Economy*, 107(1):1–39, February 1999.

Songzi Du. Robust mechanisms under common valuation. Technical report, Working Paper, Simon Fraser University, January 2018.

Florian Ederer, Richard Holden, and Margaret Meyer. Gaming and strategic opacity in incentive provision. Discussion paper, Cowles Foundation, January 2014.

Jeffrey Ely. Beeps. *American Economic Review*, 107(1):31–53, January 2017.

Jeffrey Ely, Alexander Frankel, and Emir Kamenica. Suspense and surprise. *Journal of Political Economy*, 123(1):215–260, February 2015.

Larry Epstein and Martin Schneider. Learning under ambiguity. *Review of Economic Studies*, 74(4):1275–1303, October 2007.

Mike Felgenhauer and Elisabeth Schulte. Strategic private experimentation. *American Economic Journal: Microeconomics*, 6(4):74–105, November 2014.

Alexander Frankel. Aligned delegation. *American Economic Review*, 104(1):66–83, January 2014.

Drew Fudenberg and Jean Tirole. Moral hazard and renegotiation in agency contracts. *Econometrica*, 58(6):1279–1319, November 1990.

Richard Fullerton, Bruce Linster, Michael McKee, and Stephen Slate. Using auctions to reward tournament winners: Theory and experimental investigations. *RAND Journal of Economics*, 33(1):62–84, Spring 2002.

Chishio Furukawa. Unbiased publication bias: Theory and evidence. Technical report, Massachusetts Institute of Technology, January 2017.

GAO. F-35 joint strike fighter: Continued oversight needed as program plans to begin development of new capabilities. Report to Congressional Committees 16-390, April 2016.

Daniel F. Garrett. Intertemporal price discrimination: Dynamic arrivals and changing values. *American Economic Review*, 106(11):3275–3299, November 2016.

Matthew Gentzkow and Emir Kamenica. Costly persuasion. *American Economic Review: Papers and Proceedings*, 104(5):457–462, May 2014.

Matthew Gentzkow and Emir Kamenica. A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review: Papers and Proceedings*, 106(5):597–601, May 2016.

Edward Glaeser. Researcher incentives and empirical methods. In Andrew Caplin and Andrew Schotter, editors, *The Foundations of Positives and Normative Economics*, pages 300–319. Oxford University Press, Oxford, 2008.

Jerry Green and Jean-Jacques Laffont. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456, July 1986.

Marina Halac and Pierre Yared. Commitment vs. flexibility with costly verification. Technical report, Columbia Business School, 2017.

Benjamin Handel and Kanishka Misra. Robust new product pricing. *Marketing Science*, 34(6):864–881, November 2015.

Emeric Henry. Strategic disclosure of research results: The cost of proving your honesty. *The Economic Journal*, (119):1036–1064, July 2009.

Florian Hoffmann, Roman Inderst, and Marco Ottaviani. Persuasion through selective disclosure: Implications for marketing, campaigning and privacy regulation. Working paper, Bocconi University, November 2014.

Bengt Holmstrom. Managerial incentive problems: A dynamic perspective. *Review of Economic Studies*, 66(1):169–182, January 1999.

Bengt Holmström and Paul Milgrom. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design *Journal of Law, Economics and Organization*, 7:24-52, January 1991.

Shota Ichihashi Limiting Sender's Information in Bayesian Persuasion. Working paper, Stanford University, August 2017.

John P. A. Ioannidis et al. Increasing value and reducing waste in research design, conduct and analysis. *The Lancet*, 383(9912):166–175, January 2014.

Philippe Jehiel. On transparency in organizations. *Review of Economic Studies*, (Forthcoming), 2015.

Justin P. Johnson and David P. Myatt. On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3):756–784, June 2006.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, October 2011.

Bralind Kiri, Nicola Lacetera, and Lorenzo Zirulia. Above a swamp: A theory of high-quality scientific production. Working paper, NBER, May 2015.

Anton Kolotilin. Experimental design to persuade. *Games and Economic Behavior*, 90:215–226, March 2015.

Anton Kolotilin. Experimental design to persuade. *Games and Economic Behavior*, 90:215–226, March 2015.

Anton Kolotilin, Ming Li, Tymofiy Mylovanov, and Andriy Zapechelnyuk. Persuasion of a privately informed receiver. *Econometrica*, Forthcoming, 2017.

Eugen Kovac and Tymofiy Mylovanov. Stochastic mechanisms in settings without monetary transfers. *Journal of Economic Theory*, 144(4):1373–1395, 2009.

Nicola Lacetera and Lorenzo Zirulia. The economics of scientific misconduct. *Journal of Law, Economics and Organization*, 27(3):215–260, February 2014.

Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2002.

Ying Liu. Robust pricing in discrete time. Technical report, Working Paper, NYU Stern, 2016.

Elliot Lipnowski and Doron Ravid. Cheap Talk under Transparent Motives. Working paper, University of Chicago, May 2015.

Kate Manuel. Competition in federal contracting: An overview of the legal requirements. Technical report, Congressional Research Service, 2010.

David Martimort and Aggey Semenov. Informational effects of competition and collusion in legislative politics. *Journal of Public Economics*, 92(7):1541–1563, 2008.

Edward Miguel et al. Promoting transparency in social science research. *Science*, 343(6166):30–31, January 2014.

Tymofiy Mylovanov and Andriy Zapechelnyuk. Optimal allocation with ex-post verification and limited penalties. *American Economic Review: Papers and Proceedings*, Forthcoming, 2017.

Volcker Nocke and Michael Whinston. Merger policy with merger choice. *American Economic Review*, 103(2):1006–1033, April 2013.

Benjamin A. Olken. Promises and perils of pre-analysis plans. *Journal of Economic Perspectives*, 29(3):81–98, Summer 2015.

Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, March 2014.

Di Pei. Communication with endogenous information acquisition. Working paper, Massachusetts Institute of Technology, January 2015.

Antonio Penta. Robust dynamic implementation. *Journal of Economic Theory*, 160:280–316, December 2015.

Andrea Prat. The wrong kind of transparency. *The American Economic Review*, 95(3):862–877, June 2005.

Heikki Rantakari. A simple model of project selection with strategic communication and uncertain motives. *Journal of Economic Behavior and Organization*, 103:14–42, 2014.

Frank Riedel. Optimal stopping with multiple priors. *Econometrica*, 77(3):857–908, May 2009.

John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–189, May 1983.

Anne-Katrin Roesler and Balasz Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, Forthcoming, 2017.

Michael Rothschild and Joseph Stiglitz. Increasing risk: I. a definition. *Journal of Economic Theory*, 2(3):225–243, September 1970.

Francisco Silva. If we confess our sins. Technical report, Pontificia Universidad Catolica de Chile, 2017.

Joel Sobel. Durable goods monopoly with entry of new consumers. *Econometrica*, 59(5):1455–1485, September 1991.

Nancy Stokey. Intertemporal price discrimination. *Quarterly Journal of Economics*, 93(3):355–371, August 1979.

Roland Strausz. Deterministic mechanisms and the revelation principle. *Economics Letters*, 79:333–337, June 2003.

Dezso Szalay. The economics of clear advice and extreme options. *Review of Economic Studies*, 72(4):1173–1198, October 2005.

Curtis Taylor. Digging for golden carrots: An analysis of research tournaments. *American Economic Review*, 85(4):872–890, September 1995.

Edward Wilson. Research: increasing value, reducing waste. *The Lancet*, 383(9923):1123–1124, March 2014.

Nathan Yoder. Designing incentives for academic research. Technical report, University of Wisconsin-Madison, November 2016.

Takuro Yamashita. Implementation in weakly undominated strategies: Optimality of second-price auction and posted-price mechanism. *Review of Economic Studies*, 82:1223–1246, 2015.

Appendix A

Appendix to Chapter 1

A.1 Proofs for the Main Model

We first define the transformed distribution G in cases where F need not be continuous.

DEFINITION 1'. *Given a percentile $\alpha \in (0,1]$, define $g(\alpha)$ to be the expected value of the lowest α-percentile of the distribution F. In case F is a continuous distribution, $g(\alpha) = \frac{1}{\alpha}\int_0^{F^{-1}(\alpha)} v dF(v)$. In general, g is continuous and weakly increasing.*

Let $\underline{v}$ be the minimum value in the support of F. For $\beta \in (\underline{v}, \mathbb{E}[v]]$, define $G(\beta) = \sup\{\alpha : g(\alpha) \leq \beta\}$. We extend the domain of this inverse function to $\mathbb{R}_+$ by setting $G(\beta) = 0$ for $\beta \leq \underline{v}$ and $G(\beta) = 1$ for $\beta > \mathbb{E}[v]$.[1]

We now provide proofs of the results for the main model, in the order in which they appeared.

A.1.1 Proof of Proposition 1

Given a realized price p, minimum profit occurs when there is maximum probability of signals that lead the buyer to have posterior expectation $\leq p$. First consider the information

[1] If F does not have a mass point at $\underline{v}$, $g(\alpha)$ is strictly increasing and $G(\beta)$ is its inverse function which increases continuously. If instead $F(\underline{v}) = m > 0$, then $g(\alpha) = \underline{v}$ for $\alpha \leq m$ and it is strictly increasing for $\alpha > m$. In that case $G(\beta) = 0$ for $\beta \leq \underline{v}$, after which it jumps to m and increases continuously to 1.

structure $\mathcal{I}$ that tells the buyer whether her value is in the lowest $G(p)$-percentile or above. By definition of G, the buyer's expectation is exactly p upon learning the former. This shows that, under $\mathcal{I}$, the buyer's expected value is $\leq p$ with probability $G(p)$.

Now we show that $G(p)$ cannot be improved upon. To see this, note that it is without loss of generality to consider information structures which recommend that the buyer either "buy" or "not buy". Nature chooses an information structure that minimizes the probability of "buy." By Lemma 1 in Kolotilin (2015), this minimum is achieved by a partitional information structure, namely by recommending "buy" for $v > \alpha$ and "not buy" for $v \leq \alpha$. From this, it is easy to see that the particular information structure $\mathcal{I}$ above is the worst case.

Thus, for any realized price p, the seller's minimum profit is $p(1 - G(p))$. The proposition follows from the seller optimizing over p.

A.1.2 Proof of Proposition 2

In the main text we showed that for any deterministic price path, nature can choose an information structure that holds profit down to Π^* or lower. Here we extend the argument to any randomized pricing strategy $\sigma \in \Delta(P^T)$. For clarity, the proof will be broken down into three steps.

Step 1: Cutoff values and information structure. To begin, we define a set of cutoff values. In each period t, given previous and current prices $p_1, \ldots, p_t$, a buyer who knows her value to be v prefers to buy in the current period if and only if

$$v - p_t \geq \max_{\tau \geq t+1} \mathbb{E}[\delta^{\tau-t} \cdot (v - p_\tau)] \tag{A.1}$$

where the RHS maximizes over all stopping times that stop in the future. It is easily seen that there exists a unique value v_t such that the above inequality holds if and only if $v \geq v_t$.[2]

[2]This follows by observing that both sides of the inequality are strictly increasing in v, but the LHS increases faster.

Thus, v_t is defined by the equation

$$v_t - p_t = \max_{\tau \geq t+1} \mathbb{E}[\delta^{\tau-t} \cdot (v_t - p_\tau)] \tag{A.2}$$

and it is a random variable that depends on realized prices p^t and the expected future prices $\sigma(\cdot \mid p^t)$.

Next, let us define for each $t \geq 1$

$$w_t = \min\{v_1, v_2, \ldots, v_t\} = \min\{w_{t-1}, v_t\}. \tag{A.3}$$

For notational convenience, let $w_0 = \infty$ and $w_\infty = 0$. w_t is also a random variable, and it is decreasing over time.

Consider the following information structure $\mathcal{I}$. In each period t, the buyer is told whether or not her value is in the lowest $G(w_t)$-percentile. Providing this information requires nature to know w_t, which depends only on the realized prices and the seller's (future) pricing strategy.

Step 2: Buyer behavior. The following lemma describes the buyer's optimal stopping decision in response to σ and $\mathcal{I}$:

Lemma 1': *For any pricing strategy σ, let the information structure $\mathcal{I}$ be constructed as above. Then the buyer finds it optimal to follow nature's recommendation: she buys when told her value is above the $G(w_t)$-percentile, and she waits otherwise.*

Proof of Lemma 1'. Suppose period t is the first time that the buyer learns her value is above the $G(w_t)$-percentile. Then in particular, $w_t < w_{t-1}$ which implies $w_t = v_t$ by (A.3). Given this signal, she knows that she will receive no more information in the future (because w_t decreases over time). She also knows that her value is above the $G(w_t)$-percentile, which is greater than $w_t = v_t$, the average value below that percentile. Thus from the definition of v_t, the buyer optimally buys in period t.

On the other hand, suppose that in some period t the buyer learns her value is below the $G(w_t)$-percentile. Since w_t decreases over time, this signal is Blackwell sufficient for all previous signals. By definition of G, the buyer's expected value is $w_t \leq v_t$. Thus even without additional information in the future, this buyer prefers to delay her purchase. The promise of future information does not change the result. ■

Step 3: Profit decomposition. By this lemma, the buyer with true value in the percentile range $(G(w_{t-1}), G(w_t)]$ buys in period t. Thus, the seller's expected discounted profit can be computed as

$$\Pi = \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \cdot (G(w_{t-1}) - G(w_t)) \cdot p_t\right].$$

We rely on a technical result to simplify the above expression:

Lemma 9. *Suppose $w_t = v_t \leq w_{t-1}$ in some period t. Then*

$$p_t = \mathbb{E}\left[\sum_{s=t}^{T-1} (1-\delta)\delta^{s-t} w_s + \delta^{T-t} w_T \mid p^t\right] \tag{A.4}$$

which is a discounted sum of current and expected future cutoffs.

Using Lemma 9, we can rewrite the profit as

$$\begin{aligned}
\Pi &= \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \cdot (G(w_{t-1}) - G(w_t)) \cdot \mathbb{E}\left[\sum_{s=t}^{T-1} (1-\delta)\delta^{s-t} w_s + \delta^{T-t} w_T \mid p^t\right]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \delta^{t-1} \cdot (G(w_{t-1}) - G(w_t)) \cdot \left(\sum_{s=t}^{T-1} (1-\delta)\delta^{s-t} w_s + \delta^{T-t} w_T\right)\right] \\
&= \mathbb{E}\left[\sum_{s=1}^{T-1} (1-\delta)\delta^{s-1} w_s (1 - G(w_s)) + \delta^{T-1} w_T (1 - G(w_T))\right] \\
&\leq \Pi^*.
\end{aligned} \tag{A.5}$$

The second line is by the law of iterated expectations, because w_{t-1} and w_t only depend on the realized prices p^t. The next line follows from interchanging the order of summation, and the last inequality is because $w_s(1 - G(w_s)) \leq \Pi^*$ holds for every w_s. Hence it only remains to prove Lemma 9.

Proof of Lemma 9. We assume that T is finite. The infinite-horizon result follows from an approximation by finite horizons and the Monotone Convergence Theorem, whose details we omit. We prove by induction on $T - t$, where the base case $t = T$ follows from $w_T = v_T = p_T$.

For $t < T$, from (A.2) we can find an optimal stopping time $\tau \geq t+1$ such that

$$v_t - p_t = \mathbb{E}[\delta^{\tau-t} \cdot (v_t - p_\tau)]$$

which can be rewritten as

$$p_t = \mathbb{E}[(1-\delta^{\tau-t})v_t + \delta^{\tau-t}p_\tau]. \tag{A.6}$$

We claim that in any period s with $t < s < \tau$, $v_s \geq v_t$ so that $w_s = w_t = v_t$ by (A.3); while in period τ, $v_\tau \leq v_t$ and $w_\tau = v_\tau \leq w_{\tau-1}$. In fact, if $s < \tau$, then the optimal stopping time τ suggests that the buyer with value v_t weakly prefers to wait than to buy in period s. Thus by definition of v_s, it must be true that $v_s \geq v_t$. On the other hand, in period τ the buyer with value v_t weakly prefers to buy immediately, and so $v_\tau \leq v_t$.

By these observations, if $\tau = \infty$ (meaning the buyer never buys), we have

$$(1-\delta^{\tau-t})v_t + \delta^{\tau-t}p_\tau = v_t = \sum_{s=t}^{T-1}(1-\delta)\delta^{s-t}w_s + \delta^{T-t}w_T.$$

If $\tau \leq T$, we apply inductive hypothesis to p_τ and obtain

$$(1-\delta^{\tau-t})v_t + \delta^{\tau-t}p_\tau = \sum_{s=t}^{\tau-1}(1-\delta)\delta^{s-t}w_s + \mathbb{E}\left[\sum_{s=\tau}^{T-1}(1-\delta)\delta^{s-t}w_s + \delta^{T-t}w_T \mid p^\tau\right].$$

Plugging the above two expressions into (A.6) proves the lemma. ■

A.1.3 Proof of Lemma 2

Fix a dynamic information structure $\mathcal{I}$ and an optimal stopping time τ of the buyer. Because prices are deterministic, the distribution of signal s_t in period t only depends on realized signals (but not prices). Analogously, we can think about the stopping time τ as depending only on past and current signal realizations.

As discussed in the main text, we will construct another information structure $\mathcal{I}'$ which only reveals information in the first period, and which weakly reduces the seller's profit. Consider a signal set $S = \{\bar{s}, \underline{s}\}$, corresponding to the recommendation of "buy" and "not buy", respectively. To specify the distribution of these signals conditional on v, let nature draw signals $s_1, s_2, \cdots$ according to the original information structure $\mathcal{I}$ (and conditional on

v). If, along this sequence of realized signals, the stopping time τ results in buying the object, let the buyer receive the signal $\bar{s}$ with probability $\delta^{\tau-1}$. With complementary probability and when $\tau = \infty$, let her receive the other signal $\underline{s}$. In the alternative information structure $\mathcal{I}'$, nature reveals $\bar{s}$ or $\underline{s}$ in the first period and provides no more information afterwards.

We claim that under $\mathcal{I}'$, the buyer receiving the signal $\underline{s}$ has expected value at most p_1. We actually show something stronger, namely that the buyer has expected value at most p_1 conditional on the signal $\underline{s}$ and *any* realized signal s_1.[3] To prove this, note that since stopping at time τ is weakly better than stopping at time 1, we have

$$\mathbb{E}[v \mid s_1] - p_1 \leq \mathbb{E}^{s_2,\cdots,s_T}\left[\delta^{\tau-1}(\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau] - p_\tau)\right]. \tag{A.7}$$

Here and later, the superscripts over the expectation sign highlight the random variables which the expectation is with respect to. In this case they are $s_2, \ldots, s_T$, whose distribution is governed by the original information structure $\mathcal{I}$ and the realized signal s_1.

Since $p_\tau \geq p_1$, simple algebra reduces (A.7) to the following.

$$\mathbb{E}[v \mid s_1] \leq \mathbb{E}^{s_2,\cdots,s_T}\left[\delta^{\tau-1}\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau] + (1-\delta^{\tau-1})p_1\right]. \tag{A.8}$$

Doob's Optional Sampling Theorem says that $\mathbb{E}[v \mid s_1] = \mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau]\right]$. Thus we derive the inequality:

$$p_1 \geq \frac{\mathbb{E}^{s_2,\cdots,s_T}[(1-\delta^{\tau-1}) \cdot \mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau]]}{\mathbb{E}^{s_2,\cdots,s_T}[1-\delta^{\tau-1}]}. \tag{A.9}$$

The denominator $\mathbb{E}^{s_2,\cdots,s_T}[1-\delta^{\tau-1}]$ can be rewritten as $\mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{P}(\underline{s} \mid s_1, s_2, \ldots, s_T)\right]$, which is the probability of $\underline{s}$ given s_1. Because τ is a stopping time, the numerator in (A.9) can be rewritten as

$$\mathbb{E}^{s_2,\cdots,s_T}\left[(1-\delta^{\tau-1}) \cdot \mathbb{E}[v \mid s_1, s_2, \cdots, s_T]\right]$$

[3] Technically we only consider those s_1 such that $\underline{s}$ occurs with positive probability given s_1.

which can be further rewritten as

$$\mathbb{E}^{s_2,\cdots,s_T}\left[(1-\delta^{\tau-1})\cdot\mathbb{E}[v\mid s_1,s_2,\cdots,s_T,\underline{s}]\right]$$

because $\underline{s}$ does not provide more information about v beyond $s_1,\ldots,s_T$.

With these, (A.9) states that

$$p_1 \geq \frac{\mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{P}(\underline{s}\mid s_1,s_2,\ldots,s_T)\cdot\mathbb{E}[v\mid s_1,s_2,\cdots,s_T,\underline{s}]\right]}{\mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{P}(\underline{s}\mid s_1,s_2,\ldots,s_T)\right]} = \mathbb{E}[v|s_1,\underline{s}] \tag{A.10}$$

just as we claimed.

Thus, under the information structure $\mathcal{I}'$ constructed above, a buyer who receives the signal $\underline{s}$ has expected value at most p_1, which is also less than any future price. Since information only arrives in the first period, all sale happens in the first period to the buyer with the signal $\bar{s}$. The probability of sale is at most $\mathbb{E}[\delta^{\tau-1}]$, and the seller's profit is at most $\mathbb{E}[\delta^{\tau-1}]\cdot p_1$. This is no more than $\mathbb{E}[\delta^{\tau-1}\cdot p_\tau]$, the discounted profit under the original dynamic information structure. We have thus proved that with a deterministic and non-decreasing price path, the seller's profit is at least what he would obtain by selling only once at the price p_1. Taking $p_1 = p^*$ proves the lemma.

A.1.4 Proof of Theorem 1

By the previous lemma, a constant price path p^* delivers expected un-discounted profit Π^* from each arriving buyer. This matches the upper bound given by Proposition 2 and shows that always charging p^* is optimal. Moreover, suppose p^* is unique, then from (A.5) we see that the seller's profit from the first buyer equals Π^* only if $w_s = p^*$ almost surely. This together with Lemma 9 implies $p_1 = p^*$ almost surely. Analogous argument for later buyer shows that the seller must always charge p^* to achieve the maxmin profit. Hence the proposition.

A.2 Proofs for the Alternative Timing Model

A.2.1 Proof of Theorem 2

Throughout, we represent the robust selling mechanism in Du (2018) by a random price, with c.d.f. $D(x)$; the details of this distribution can be found later in (A.24), but they are not relevant for this proof. Because nature can provide each arriving buyer with the Roesler-Szentes information structure (1.3), the seller at most obtains Π_{RSD} from each buyer. To complete the proof, we will construct a dynamic pricing strategy that yields Π_{RSD} from each buyer.

The following lemma proves the outcome-equivalence between static and dynamic pricing strategies, and it may be of independent interest:

Lemma 10. *Fix any continuous distribution function D, any horizon T and any discount factor $\delta \in (0,1)$. There exists a distribution of prices $\sigma \in \Delta(p^T)$ such that if a buyer arrives in period t and knows her value to be v, then her discounted probability of purchasing the object (discounted to period t) is equal to $D(v)$.*

In words, for any static pricing strategy there is a dynamic pricing strategy which does not condition on buyers' arrival times, but which results in the same outcome as the static prices for every type of each arriving buyer.

We state the lemma for continuous distributions so that the buyer's optimal stopping time is almost surely unique. From Du (2018), Du's distribution D is continuous except when it is a point-mass on W. In the latter case $\Pi_{RSD} = \Pi^*$, and Theorem 2 follows from Theorem 1.

Lemma 10 is useful for our problem because it implies, via the Revenue Equivalence Theorem, that a seller using strategy σ obtains the same profit from any buyer as if he sells only once to this buyer at a random price distributed according to D. This is true whenever the buyer's value distribution is determined upon arrival and fixed over time, which is what we assume for the current proposition. Since Du's static mechanism guarantees profit Π_{RSD} from every buyer, the proposition will follow once we prove the lemma.

Proof of Lemma 10. We will first prove the result for $T = 2$, then generalize to all finite T and lastly discuss $T = \infty$.

Step 1: The case of two periods. In the second period, regardless of realized p_1 the seller should charge a random price drawn from D. This achieves the desired allocation probabilities for the second buyer.

Consider the first buyer. For any price p_1 in the first period, define v_1 as the cutoff indifferent between buying at price p_1 or waiting till the next period and facing the random price drawn from D. That is,

$$v_1 - p_1 = \delta \cdot \mathbb{E}^{p_2 \sim D}\left[\max\{v_1 - p_2, 0\}\right]. \tag{A.11}$$

As p_1 varies according to the seller's pricing strategy σ, v_1 is a random variable. As in the proof of Proposition 2, we define $w_1 = v_1$ and $w_2 = \min\{v_1, p_2\}$, where p_2 is independently drawn according to D.

If the buyer has value $x \geq w_1$, she buys in the first period. Otherwise if she has value $w_1 > x \geq w_2$, she buys in the second period. The discounted purchasing probability of such a buyer is thus

$$\mathbb{P}^{w_1}[x \geq w_1] + \delta \cdot \mathbb{P}^{w_1, w_2}[w_1 > x \geq w_2] = (1-\delta) \cdot \mathbb{P}^{w_1}[x \geq w_1] + \delta \cdot \mathbb{P}^{w_2}[x \geq w_2].$$

Let w be the random variable that satisfies $w = w_1$ (or w_2) with probability $1-\delta$ (or δ), then the seller seeks to ensure that w is distributed according to D.

Suppose H is the c.d.f. of v_1. Since $w_1 = v_1$ and $w_2 = \min\{v_1, p_2\}$, the probability that w is *greater than* x is given by $(1-\delta)(1-H(x)) + \delta(1-H(x))(1-D(x))$.[4] This has to be equal to $1 - D(x)$, which implies

$$1 - H(x) = \frac{1 - D(x)}{1 - \delta D(x)}. \tag{A.12}$$

We are left with the task of finding a first-period price distribution under which $v_1 \sim H$. This can be done because the random variables v_1 and p_1 are in a one-to-one relation (see

[4] $1 - H(x)$ is the probability that $w_1 > x$, and $(1 - H(x))(1 - D(x))$ is the probability that $w_2 > x$.

(A.11)). We have proved the lemma for $T = 2$.

Before proceeding, we remark that (A.12) implies the distribution H has the same support as D. However, (A.11) suggests that when v_1 achieves the maximum of this support, p_1 is in general strictly smaller than v_1 (unless the support is a singleton point, a case we have discussed). Intuitively, charging this maximum price in the first period leads to delayed purchase by buyers with high values, which is costly for the seller. On the other hand, the minimum price p_1 is indeed equal to the minimum of the support of D, which we denote by W; when D is Du's distribution, this is the same W as in the Roesler-Szentes information structure (1.3).

Step 2: Extension to finite T. We conjecture a pricing strategy σ that is independent across periods: $d\sigma(p_1, \ldots, p_T) = d\sigma_1(p_1) \times \cdots \times d\sigma_T(p_T)$, where we interpret each σ_t as a distribution. Define the cutoff values $v_1, \ldots, v_T$ as in (A.2). Note that due to independence, v_t only depends on current price p_t but not on previous prices.

Consider a buyer who arrives in period t. We can generalize the previous arguments and show that if she knows her value to be x, then her discounted purchasing probability is $\mathbb{P}[w^{(t)} \leq x]$. The random variable $w^{(t)}$ is described as follows: for $t \leq s \leq T-1$, $w^{(t)} = \min\{v_t, v_{t+1}, \ldots, v_s\}$ with probability $(1-\delta)\delta^{s-t}$; and with remaining probability δ^{T-t}, $w^{(t)} = \min\{v_t, v_{t+1}, \ldots, v_T\}$.

The result of the lemma requires each $w^{(t)}$ to be distributed according to D. Simple calculation shows this is the case if $v_T \sim D$ and $v_1, \ldots, v_{T-1} \sim H$ (since v_t depends only on p_t, they are independent random variables).[5] We can then solve for the price distributions $\sigma_1, \ldots, \sigma_T$ by backward induction: σ_T must be D, and once the prices in period $t+1, \ldots, T$ are determined, there is a one-to-one relation between p_t and v_t by (A.2). Thus, the distribution of p_t is uniquely pinned down by the desired distribution of v_t.

Step 3: The infinite horizon case. If $T = \infty$, we look for price distributions $\sigma_1, \sigma_2, \ldots$ such

[5]The reason $H(x)$ should be the c.d.f. of v_1 is best understood in the infinite horizon problem (see below). Under stationarity, the buyer with value x buys in period t with probability $H(x)$, conditional on not buying previously. Thus the discounted allocation probability is $\sum_t \delta^{t-1}(1-H(x))^{t-1}H(x)$. Setting this equal to $D(x)$ yields (A.12).

that $v_1, v_2, \cdots \sim H$. We conjecture a stationary σ_t. Recall that the cutoff v_1 is defined by

$$v_1 - p_1 = \max_{\tau \geq 2} \mathbb{E}\left[\delta^{\tau-1}(v_1 - p_\tau)\right]. \tag{A.13}$$

The stopping problem on the RHS is stationary. Thus when $p_2 < p_1$ the buyer stops in period 2 and receives $v_1 - p_2$; otherwise she continues and receives $v_1 - p_1$. (A.13) thus reduces to

$$v_1 - p_1 = \delta \cdot \mathbb{E}^{p_2}\left[\max\{v_1 - p_1, v_1 - p_2\}\right]$$

which can be further simplified to

$$v_1 = p_1 + \frac{\delta}{1-\delta} \cdot \mathbb{E}^{p_2}\left[\max\{p_1 - p_2, 0\}\right]. \tag{A.14}$$

Let $P(x)$ denote the c.d.f. of p_1 (and of p_2). When $p_1 = x$, (A.14) implies

$$v_1 = x + \frac{\delta}{1-\delta} \cdot \int_0^x (x - z)\, dP(z) = x + \frac{\delta}{1-\delta} \int_0^x P(z)\, dz.$$

Thus v_1 has c.d.f. $H(x)$ if and only if

$$P(x) = H\left(x + \frac{\delta}{1-\delta} \int_0^x P(z)\, dz\right). \tag{A.15}$$

To solve for $P(x)$, we let

$$Q(x) = x + \frac{\delta}{1-\delta} \int_0^x P(z)\, dz; \qquad U(y) = 1 + \frac{\delta}{1-\delta} H(y) = \frac{1}{1 - \delta D(y)}. \tag{A.16}$$

(A.15) is the differential equation

$$U(Q(x)) = Q'(x). \tag{A.17}$$

Put $V(y) = \int_0^y (1 - \delta D(z))\, dz$, so that $V'(y) = \frac{1}{U(y)}$. Then

$$\frac{\partial V(Q(x))}{\partial x} = V'(Q(x)) \cdot Q'(x) = \frac{Q'(x)}{U(Q(x))} = 1. \tag{A.18}$$

Inspired by the analysis for finite T, we conjecture that the minimum value of p_1 is W. That is, we conjecture $Q(W) = W$. Since $V(W) = W$, we deduce from (A.18) that $Q(x)$ is characterized by

$$V(Q(x)) = x \ \text{ with } \ V(y) = \int_0^y (1 - \delta D(z))\, dz. \tag{A.19}$$

Since V is strictly increasing, there is a unique solution $Q(x)$ to the above equation, and the corresponding distribution of prices is

$$P(x) = \frac{1-\delta}{\delta} \cdot (Q'(x) - 1). \tag{A.20}$$

Lemma 10 is proved, and so is Theorem 2. ■

A.2.2 Proof of Theorem 3

Consider a constant price p randomly drawn according to the Du distribution. Recall that p is supported on $[S, W]$, and its density is $\frac{1}{\log \frac{S}{W}} \cdot \frac{1}{p}$. We show the seller's discounted expected profit is at least W.

By assumption, each buyer's expected value follows a martingale process $v_1, v_2, \ldots$ that is autonomous (independent of the realized p). As mentioned in the main text, we define a sequence of cutoff prices adapted to the v-process:

$$v_t - r_t = \max_{\tau > t} \mathbb{E}[\delta^{\tau - t}(v_\tau - r_t)]$$

and then

$$q_t = \max \{r_1, \ldots, r_t\}.$$

This is exactly dual to the definition of cutoff values, and whenever $q_t = r_t \geq q_{t-1}$, we have (See Lemma 9):

$$v_t = \mathbb{E}\left[\sum_{s \geq t} (1-\delta)\delta^{s-t} q_s \mid v_1, \ldots, v_t\right].$$

If the random price p satisfies $q_{t-1} \leq p < q_t$, then purchase occurs in period t. Total profit is thus:

$$
\begin{aligned}
\Pi &= \mathbb{E}\left[\sum_{t\geq 1}\delta^{t-1}\int_{q_{t-1}}^{q_t} p\, dD(p)\right] \\
&= \frac{1}{\log\frac{S}{W}}\cdot\mathbb{E}\left[\sum_{t\geq 1}\delta^{t-1}(\pi(q_t)-\pi(q_{t-1}))\right] \\
&= \frac{1}{\log\frac{S}{W}}\cdot\mathbb{E}\left[\sum_{t\geq 1}(1-\delta)\delta^{t-1}\pi(q_t)\right]
\end{aligned}
$$

where we define $\pi(y)=\min\{(y-W)^+, S-W\}$ to be the integral of $\log\frac{S}{W}\cdot p\, dD(p)$, and use the convention that $\pi(q_0)=0$. In other words, $\pi(y)=0$ for $y\leq W$, $\pi(y)=y-W$ for $y\in[W,S]$ and $\pi(y)=S-W$ for $y\geq S$. Define:

$$
\hat{\pi}(y)=\min\{y-W, S-W\}=y-w-(y-S)^+
$$

to be a modified version of the function π.[6]

Indeed, $\hat{\pi}$ is smaller than π and strictly so when $y\leq W$. Then we have

$$
\begin{aligned}
\log\frac{S}{W}\cdot\Pi &= \mathbb{E}\left[\sum_{t\geq 1}(1-\delta)\delta^{t-1}\pi(q_t)\right] \\
&\geq \mathbb{E}\left[(1-\delta)\delta^{t-1}\hat{\pi}(q_t)\right] \\
&= \mathbb{E}\left[(1-\delta)\delta^{t-1}(q_t-W-(q_t-S)^+)\right] \\
&= v_0-W-\mathbb{E}\left[(1-\delta)\delta^{t-1}(q_t-S)^+)\right]
\end{aligned}
$$

where we use the fact that the ex-ante expected value v_0 is a discounted sum of cutoff prices.

Let γ be a stopping time adapted to the v-process such that q_γ first exceeds S. Then we

[6]Intuitively, π coincides with $\hat{\pi}$ if the threshold prices never fall below W. We cannot assume this a priori, although it is natural to expect that doing so would not be worst case, just as inducing a belief lower than the price to a non-purchasing buyer is not worst case in the main model.

can continue the above computation as follows

$$
\begin{aligned}
\log \frac{S}{W} \cdot \Pi &\geq v_0 - W - \mathbb{E}\left[(1-\delta)\delta^{t-1}(q_t - S)^+\right] \\
&= v_0 - W - \mathbb{E}\left[\delta^{\gamma-1} \sum_{t \geq \gamma} (1-\delta)\delta^{t-\gamma}(q_t - S)\right] \\
&= v_0 - W - \mathbb{E}\left[\delta^{\gamma-1}(v_\gamma - S)\right] \\
&\geq v_0 - W - \mathbb{E}[(v_\infty - S)^+] \\
&= \mathbb{E}[v_\infty - W - (v_\infty - S)^+] \\
&= \mathbb{E}[\hat{\pi}(v_\infty)].
\end{aligned}
$$

The inequality holds since if γ is finite, then $v_\gamma - S \leq \mathbb{E}[(v_\infty - S)^+ \mid v_1, \ldots, v_\gamma]$ by convexity. And if γ is infinite, then $\delta^{\gamma-1}(v_\gamma - S) = 0 \leq (v_\infty - S)^+$.

To summarize, we first showed $\log \frac{S}{W} \cdot \Pi \geq \mathbb{E}\left[(1-\delta)\delta^{t-1}\hat{\pi}(q_t)\right]$, and since $\hat{\pi}$ is *concave*, this is smaller than $\hat{\pi}(\mathbb{E}[(1-\delta)\delta^{t-1}q_t]) = \hat{\pi}(v_0)$. However, the lower bound $\mathbb{E}[\hat{\pi}(v_\infty)]$ holds by concavity, essentially because the distribution of v_∞ is more spread out than the cutoff prices q_t.

Letting $\tilde{F}$ denoting the distribution of v_∞, we have:[7]

$$
\begin{aligned}
\log \frac{S}{W} \cdot \Pi &\geq \mathbb{E}[\hat{\pi}(v_\infty)] \\
&= \int_0^S (v - W) d\tilde{F}(v) + (S - W)(1 - \tilde{F}(S)) \\
&= \tilde{F}(S)(S - W) - \int_0^S \tilde{F}(v) dv + (S - W)(1 - \tilde{F}(S)) \\
&= S - W - \int_0^S \tilde{F}(v) dv \\
&\geq S - W - \int_0^S F(v) dv \\
&= \log \frac{S}{W} \cdot W
\end{aligned}
$$

The first inequality follows from F being a mean preserving spread of $\tilde{F}$, and the second

[7]The stronger result $\log \frac{S}{W} \cdot \Pi \geq \mathbb{E}[\pi(v_\infty)]$ would mean that profit is minimized by revealing all information at once, which would easily complete the proof. But in order to use concavity, we have had to work with $\hat{\pi}$ rather than π.

follows from (1.3). Hence $\Pi \geq W$ as desired. ■

A.2.3 Proof of Claim 4

The proof is somewhat long, and we will present it in several steps. First, we review some properties of Du's static mechanism. Next, we focus on the pricing strategy σ^D that we constructed in the preceding proof. We construct a *dynamic* information structure (for the first buyer) that yields profit below Π_{RSD}. This proves the proposition assuming that the seller uses the strategy σ^D. Lastly, we apply continuity arguments and extend the result to any pricing strategy σ.

Step 1: Properties of the Optimal Single-Period Mechanism. While the solution to the model with a single-period was provided by Roesler and Szentes (2017) and Du (2018), our proof of Claim 4 requires us to demonstrate some properties of the solution. We defer this to Appendix A.4.1, since they are orthogonal to information arrival dynamics.

Lemma 11. *For generic distributions F, there is a unique random-price mechanism that achieves Π_{RSD}.*

For reference, the random-price mechanism which achieves the optimum Π_{RSD} is given by:

$$D(x) = \begin{cases} 0 & x \in [0, W) \\ \frac{\log \frac{x}{W}}{\log \frac{S}{W}} & x \in [W, S) \\ 1 & x \in [S, 1] \end{cases}, \tag{A.21}$$

where $S \in (W, B]$ is characterized by

$$\int_0^S F_W^B(v)\, dv = \int_0^S F(v)\, dv, \tag{A.22}$$

where F_W^B is the Roesler-Szentes worst-case information structure (1.3).

Step 2: The Information Structure. Consider now the model with two periods and one buyer arriving in each period. The problem for the second buyer is static, so nature can choose an information structure that yields profit at most Π_{RSD}.

We construct the following dynamic information structure $\mathcal{I}$ for the *first* buyer:

- In the first period, nature provides the Roesler-Szentes information structure. We denote the buyer's unbiased signal by $\tilde{v}$ (which is also her posterior expected value), so as to distinguish from her true value v. Note that $\tilde{v} \sim F_W^B$.

- In the second period, given the realized price p_1 as well as the buyer's expected value $\tilde{v}$ in the first period, nature reveals the buyer's true value v if and only if $\tilde{v} \geq v_1(p_1)$. Otherwise nature provides no additional information. Here the cutoff $v_1(p_1)$ is defined as usual, assuming no information arrives in the second period:

$$v_1 - p_1 = \delta \cdot \mathbb{E}^{p_2 \sim \sigma(\cdot|p_1)} \left[\max\{v_1 - p_2, 0\}\right].$$

  Note that in general, the distribution of p_2 may depend on p_1.

Intuitively, nature targets the buyer who prefers to buy in the first period when she does not expect to receive information in the second period. By promising full information to such a buyer in the future, nature potentially delays her purchase and reduces the seller's profit. In what follows we formalize this intuition.

Step 3: Buyer behavior and seller profit. To facilitate the discussion, we consider another information structure $\mathcal{I}'$ in which nature reveals $\tilde{v}$ in the first period but does nothing in the second period. Under $\mathcal{I}'$, the buyer's value distribution F_W^B does not change over time. Thus by Stokey (1979), the seller's profit would at most be Π_{RSD}. We will show that the seller's profit under the *dynamic* information structure $\mathcal{I}$ could only be lower than under $\mathcal{I}'$ (for any pricing strategy), and we also characterize when the comparison is strict.

There are three possibilities: first, if the price p_1 is relatively high so that $\tilde{v} < v_1(p_1)$, then the buyer does not buy in the first period under $\mathcal{I}'$. This is also her optimal decision under $\mathcal{I}$, because she will not receive extra information in the second period. Secondly, if the price is very low, then under both $\mathcal{I}$ and $\mathcal{I}'$ the buyer buys in the first period. Lastly, for some intermediate prices the buyer buys in the first period under $\mathcal{I}'$ but not under $\mathcal{I}$; the opposite situation cannot occur because $\mathcal{I}$ provides more information than $\mathcal{I}'$ in the second period, and the buyer's incentive to wait could only be stronger.

Thus, when nature provides $\mathcal{I}$ rather than $\mathcal{I}'$, the seller's profit changes only in the last possibility above. Let us show that whenever the buyer delays her purchase from the first period to the second, the seller's profit decreases by at least $(1-\delta)W$. This is because when the buyer chooses to not buy in the first period, the discounted social surplus decreases by at least $(1-\delta)\tilde{v}$. Since the buyer's payoff cannot decrease (because she *chooses* to delay purchase), the loss must come from the seller's discounted profit.

To summarize, we have shown:

Lemma 12. *Consider the information structures $\mathcal{I}$ and $\mathcal{I}'$ constructed above. The seller's profit under $\mathcal{I}'$ is no greater than Π_{RSD}, and his profit under $\mathcal{I}$ is at least smaller by $(1-\delta)W$ times the probability that the buyer delays purchase.*

Step 4: Proof of the claim for σ^D. Let σ^D be the pricing strategy given by Lemma 10, which we recall is robust to information that arrives only once (for each buyer). Here we show that under the dynamic information structure $\mathcal{I}$, the seller's profit from the first buyer is strictly less than Π_{RSD}.

Recall from the proof of Lemma 10 that under σ^D, the price in the second period p_2 is drawn from Du's distribution D, independent of p_1. On the other hand, p_1 is (continuously supported) on a smaller interval $[W, S_1]$, with $W < S_1 < S$; more precisely, the distribution of p_1 is determined by the condition that $v_1(p_1) \sim H$ (see (A.12)).

Suppose the buyer receives unbiased signal $\tilde{v} \in (W, S)$ in the first period. She delays her purchase at some price $p_1 \in (W, S_1)$ under information structure $\mathcal{I}$ (compared to $\mathcal{I}'$) if and only if knowing her true value *strictly* improves her expected utility in the second period; because $p_2 \sim D$ regardless of p_1, delay occurs if p_1 is smaller than but close to $v_1^{-1}(\tilde{v})$. We will demonstrate a positive measure of such $\tilde{v}$, so that the buyer delays purchase with strictly positive probability.

Now recall from Step 1 that a signal $\tilde{v} < S$ is only received when the true value also satisfies $v < S$. Because we assume $\Pi_{RSD} > \Pi^*$, Proposition 11 gives $W > \underline{v}$. Thus a positive measure of signals $\tilde{v} \in (W, S)$ is received when the true value v belongs to the interval $[\underline{v}, W)$.[8] We claim that for any such $\tilde{v}$, knowing the true value in the second period strictly benefits the buyer. This is because according to her expected value $\tilde{v} > W$, the buyer in the second period buys at some price p_2; but if she were informed that $v < W$, she would not buy at any price p_2 (which is at least W). This proves that by providing $\mathcal{I}$ rather than $\mathcal{I}'$, nature induces a positive probability of delay. By Lemma 12, we deduce that profit from the first buyer is less than Π_{RSD}.

Step 5: Proof for an arbitrary pricing strategy σ. Finally, we turn to prove the proposition in its full generality. The argument is as follows (omitting technical details): suppose for contradiction that some pricing strategy σ guarantees profit almost Π_{RSD} from each buyer. Then because D is uniquely optimal in the one-period problem, the distribution of p_2 conditional on p_1 is "close" to D (in the Prokhorov metric) with high probability; otherwise nature could sufficiently damage the seller's profit from the second buyer. Next, we can similarly show that the distribution of $v_1(p_1)$ under σ is close to H, which is its distribution if $\sigma = \sigma^D$.[9] The rest of the proof proceeds as in Step 4: a positive measure of signals $\tilde{v} \in (W, S)$ is received when the true value satisfies $v < W$. For such $\tilde{v}$, full information in the second period is strictly valuable, and the buyer delays purchase if $v_1(p_1)$ is smaller than but close to $\tilde{v}$. By what we have shown, this occurs with strictly positive probability. But then Lemma 12 implies profit from the first buyer is bounded away from Π_{RSD} under $\mathcal{I}$, leading to a contradiction. The proof of Claim 4 is complete. ■

Let us conclude by commenting on the assumption that Du's mechanism is uniquely

[8] If $\Pi_{RSD} = \Pi^*$, then Proposition 11 implies $W = \underline{v} = p^*$ and Du's distribution is a mass-point at W. Information in the second period is irrelevant, because a buyer waiting till the second period always buys at price $p_2 = W = \underline{v}$.

[9] Consider nature choosing $\tilde{F}$ in the first period and doing nothing afterwards. The seller's profit from the first buyer can be written as $\mathbb{E}^w[w(1 - \tilde{F}(w)]$, where the random variable w equals $v_1(p_1)$ with probability $1 - \delta$ and it equals $\min\{v_1(p_1), p_2\}$ with probability δ (see (A.5)). The distribution of w must be close to D, otherwise nature could choose $\tilde{F}$ and damage profit from the first buyer. Since p_2 is approximately distributed according to D, we can derive as in the proof of Lemma 10 that $v_1(p_1)$ must be approximately distributed according to H.

optimal. Suppose this assumption fails, so that another point $\hat{S} > S$ satisfies (A.25). This means there are two different Du distributions D and $\hat{D}$, supported on $[W, S]$ and $[W, \hat{S}]$ respectively. On their supports, both of these distributions have density proportional to $\frac{1}{p}$ (see (A.24)). This observation allows us to write

$$\hat{D} = \alpha D + (1 - \alpha) E \tag{A.23}$$

with $\alpha \in (0, 1)$ is a scalar and E is a distribution supported on $[S, \hat{S}]$ (again with density proportional to $\frac{1}{p}$).

When such non-uniqueness occurs, the previous proof of Claim 4 fails. Specifically, in Step 5, we are not able to deduce that σ is "close" to either σ^D or $\sigma^{\hat{D}}$. In fact, the following pricing strategy σ guarantees profit Π_{RSD} from the second buyer as well as from the first buyer, if nature chooses the information structure $\mathcal{I}$ in Step 2.

- The seller chooses a distribution of p_1 so that $v_1(p_1) \sim E$, which is supported on $[S, S']$. Here $v_1(p_1)$ is defined by the usual indifference condition $v_1 - p_1 = \delta \cdot \mathbb{E}^{p_2 \sim D}[\max\{v_1 - p_2, 0\}]$.
- Independent of the realized p_1, the seller draws $p_2 \sim D$, supported on $[W, S]$.

Because the price in the second period follows a Du distribution, the seller's profit from the second buyer is at least Π_{RSD}. For the first buyer, consider first the information structure $\mathcal{I}'$ as in Step 3, where nature reveals $\tilde{v} \sim F_W^B$ in the first period and no additional information afterwards. As shown in Footnote 9, the seller's profit from this buyer is $\mathbb{E}^w[w(1 - F_W^B(w)]$. This is as in the one-period model, where the seller charges price w and nature provides the Roesler-Szentes information structure.

Recall that w is a random variable that equals $v_1(p_1)$ with probability $1 - \delta$ and $\min\{v_1(p_1), p_2\}$ with complementary probability. Because $v_1(p_1) \sim E$, whose support is strictly above the support of p_2, we deduce that $w \sim \delta D + (1 - \delta) E$. Thus, by (A.23), the distribution of w is a convex combination of D and $\hat{D}$ whenever $\delta \geq \alpha$. Since the seller ensures profit Π_{RSD} by using a random price distributed according to either D or $\hat{D}$, he does just as well by charging w. We have thus shown that profit from the first buyer is at least Π_{RSD} under information structure $\mathcal{I}'$.

Moreover, we claim that when nature provides $\mathcal{I}$ rather than $\mathcal{I}'$, no buyer delays her purchase. To see this, consider a buyer who purchases in the first period under $\mathcal{I}'$. By definition of v_1 and the fact that $v_1 \sim E$, this means the buyer's signal $\tilde{v}$ in the first period satisfies $\tilde{v} \geq v_1(p_1) \geq S$. But then her true value v must also be at least S, as we showed in Step 1. Such a buyer purchases at any price $p_2 \in [W, S]$ regardless of any information in the second period. Thus, although nature promises future information under $\mathcal{I}$, this information does not improve the buyer's expected utility in the second period. Consequently the buyer's behavior under $\mathcal{I}$ is the same as under $\mathcal{I}'$, and profit under $\mathcal{I}$ is also equal to Π_{RSD}.

To summarize, we have constructed a pricing strategy σ such that if nature chooses the particular information structure $\mathcal{I}$ (for the first buyer), the seller's total profit is at least $(1+\delta)\Pi_{RSD}$. This explains why *our proof of* Claim 4 requires the assumption that Du's mechanism is unique. We do not know whether the result generally holds without this assumption.[10]

A.3 Proofs for the Informational Externalities Model

The Appendix presents a proof of Lemma 13 below, exhibiting the worst-case information structure in the case of short-lived buyers (i.e., buyers who can only buy upon arrival or never). The reader may find this proof help since with informational externalities, the no-delay worst case information structure does not follow from the $T = 1$ solution, in contrast to the baseline model:

Lemma 13. *Consider the model with common values and public signals suppose buyers are short lived. Fix a weakly increasing price path $(p_1, \ldots, p_T)$ with $p_1 \leq p_2 \leq \cdots \leq p_T$. Then the worst case profit is achievable by an information structure that invovles a single signal that is observed by*

[10]In other words, suppose S is not unique and suppose the seller uses the strategy σ constructed just now. We do not know whether nature can damage the seller's profit to be strictly lower than Π_{RSD} by choosing an information structure different from the $\mathcal{I}$ in our proof.

all buyers.

Proof of Lemma 3. We prove the theorem for the case of finite T by induction; the infinite horizon case follows from an approximation argument whose details are omitted. In addition to the conclusion of the Lemma, we will show that any worst case information structure induces a buyer belief lower than p_1 with probability 0. The conclusion of the Lemma holds for the case of $T = 1$ immediately, and the additional statement holds by Proposition 1.

Fix some arbitrary finite T. Consider an arbitrary dynamic information structure for the first buyer $\mathcal{I}_1$, which we assume contains at least as much information as is generated by the information structure after $t \geq 2$ (i.e., this buyer would see all information provided to later buyers after delaying). By the inductive step, we take the signal space for all buyers with arrival $a \geq 2$ to be a time s (possibly equal to $T + 1$), with the interpretation that a buyer arriving at a purchases whenever $a \leq s$, and does not purchase whenever $a > s$; this signal is commonly observed by all buyers with arrival time $a \geq 2$ (though the first buyer may not observe this).

Following Lemma 2, consider the following replacement information structure for the first buyer: let the signal space be $\{\underline{s}, 2, \ldots, T + 1\}$, and give the recommendation $\underline{s}$ if the information structure dictates that this buyer not buy, as well as with probability $1 - \delta^{\tau-1}$ if the information structure results in purchase at time τ. If the first buyer does not buy, then reveal no information to later buyers. If the first buyer does buy, then also reveal to the first buyer the time s at which purchasing will stop (recalling that all buyers with arrival time $a \geq 2$ commonly observed this signal). Note that the replacement is a public information structure.

By repeating the proof of Lemma 2, we can show that this replacement is obedient, and that the payoff from the first buyer does not increase. Furthermore, if the first buyer does not buy, then neither do any subsequent buyers, meaning that the seller's surplus is also minimized in this continuation history. On the other hand, if the first buyer does buy, then the original information structure revealed some time $s \geq 2$ at which point buyers stopped purchasing. Note that in any worst case information structure, beliefs are never lower than $p_s \geq p_1$, since otherwise a recommendation to not buy could be induced with higher probability. Hence the first buyer would still be willing to follow the recommendation in the replacement. Therefore, the replacement can only hurt the seller. ■

Proof of Theorem 5. Note first that $\lim_{\delta \to 1} \ (1-\delta)\Pi^C \leq \Pi_{RSD}$, since nature can choose F_{RS} and reveal it publicly at time 1. In general, when nature chooses distribution D, we write the seller's average profit as:

$$\Pi^C((p_t)_{t=1}^{\infty}) = \sum_{t=1}^{\infty} p_t(1 - D(p_t))\delta^{t-1}(1-\delta).$$

Let $\tilde{F}$ denote the random price distribution (with density $\tilde{f}$) corresponding to Du's mechanism (See (A.24)). Let $(p_t^\delta)_{t=1}^{\infty}$ denote a sequence of price paths such that:

$$\tilde{F}(p_t^\delta) - \tilde{F}(p_{t-1}^\delta) \to \delta^{t-1}(1-\delta)$$

with $p_1 = \Pi_{RSD} = W$. For any such price path, we have:

$$\Pi^C((p_t^\delta)_{t=1}^{\infty}) \to \int_W^B p(1-D(p))\tilde{f}(p)dp = \Pi_{RSD},$$

as claimed. ■

A.4 Additional Appendices

A.4.1 Additional Proofs for the Timing Section

Properties of Du's mechanism used in the proof of Claim 4.

For the one-period model, Du (2018) constructs a mechanism that guarantees profit Π_{RSD} regardless of the buyer's information structure. By considering the profile of interim allocation probabilities as a c.d.f., we can equivalently implement Du's mechanism as a random price with the following distribution:

$$D(x) = \begin{cases} 0 & x \in [0, W) \\ \frac{\log \frac{x}{W}}{\log \frac{S}{W}} & x \in [W, S) \\ 1 & x \in [S, 1] \end{cases} \tag{A.24}$$

Here $S \in (W, B]$ is characterized by[11]

$$\int_0^S F_W^B(v)\, dv = \int_0^S F(v)\, dv \tag{A.25}$$

where F_W^B is the Roesler-Szentes worst-case information structure (1.3). To explain further, Roesler and Szentes (2017) observe that the LHS in (A.25) must not exceed the RHS (for all S) because F is a mean-preserving spread of F_W^B. However, when W is smallest possible, this constraint must bind at some S.

The following observations will be crucial. Since the constraint $\int_0^x F_W^B(v)\, dv \leq \int_0^x F(v)\, dv$ binds at $x = S$, the first order condition gives $F_W^B(S) = F(S)$. This implies that not only F is a mean-preserving spread of F_W^B, but in fact the truncated distribution of F conditional on $v \leq S$ is also a mean-preserving spread of the corresponding truncation of F_W^B. In other words, the Roesler-Szentes information structure has the property that a buyer with true value $v \leq S$ only receives signal $\leq S$ (i.e., her posterior expected value is at most S), while a buyer with true value $v > S$ expects her value to be greater than S.

For completeness, we include a quick proof that the random price $p \sim D$ guarantees profit $W = \Pi_{RSD}$. Consider the one-period model in which nature chooses a distribution $\tilde{F}$ of the buyer's posterior expected values. Then the seller's profit is

$$\begin{aligned} \Pi = \int_W^S p(1 - \tilde{F}(p))\, dD(p) = \frac{1}{\log \frac{S}{W}} \int_W^S (1 - \tilde{F}(p))\, dp \geq \frac{1}{\log \frac{S}{W}} \left(S - W - \int_0^S \tilde{F}(p)\, dp \right) \\ \geq \frac{1}{\log \frac{S}{W}} \left(S - W - \int_0^S F(p)\, dp \right) = \frac{1}{\log \frac{S}{W}} \left(S - W - \int_0^S F_W^B(p)\, dp \right) = W. \end{aligned}$$

The penultimate equality uses (A.25) and the last one uses (1.3).

[11] S is strictly greater than W because otherwise D is a mass-point at W and $\Pi_{RSD} = \Pi^*$, contradicting the assumption of the proposition.

We note that in general, there could be more than one point S for which (A.25) holds. Thus, the maxmin optimal mechanism in one period need not be unique even if we restrict attention to the class of exponential mechanisms considered by Du (2018). But we do have the following result:

Lemma 14. *There is a unique maxmin optimal mechanism in the one-period simultaneous-move model if and only if (A.25) holds at a unique point S.*

We mention that for generic distributions F, there is a unique S that satisfies (A.25). However, the proof is tangential to the paper and we will leave it out. A sufficient condition is that $F(x)$ is convex, for example when F is uniform.[12]

Proof of Lemma 14. "Only if" is obvious, so we focus on the "if" direction. Suppose S is unique, we need to show any random price that guarantees W must follow Du's distribution D. Suppose $r(p)$ is the p.d.f. of the random price, then profit is

$$\Pi = \int_0^1 p \cdot r(p) \cdot (1 - \tilde{F}(p))\, dp. \tag{A.26}$$

Given $r(p)$, Nature's problem is to choose a c.d.f. $\tilde{F}$ to minimize Π, subject to $\int_0^x \tilde{F}(v)\, dv \leq \int_0^x F(v)\, dv$ for all $x \in (0,1]$, with equality at $x = 1$ (so that $\tilde{F}$ has the same mean as F).

By Roesler and Szentes (2017), $\tilde{F} = F_W^B$ is a solution to nature's problem. For this solution, the integral inequality constraint only binds at $x = S$. Standard perturbation techniques in the calculus of variations thus imply that $\tilde{F} = F_W^B$ cannot be improved upon only if $p \cdot r(p)$ is a constant for $p \in (W, S)$.[13] Similarly, $p \cdot r(p)$ must also be a constant on the interval $p \in (S, B)$; in fact, we can show this constant is zero.[14]

---

[12]Recall that $F(S) = F_W^B(S)$. However, $F(x) - F_W^B(x) = F(x) + \frac{W}{x} - 1$ is convex, and so it has at most two roots $x_0 < x_1$. Because $F(x) > F_W^B(x)$ for $x < x_0$, S being x_0 would contradict (A.25). Thus $S = x_1$ is unique.

[13]Suppose to the contrary that $p \cdot r(p) > p' \cdot r(p')$ for some $p, p' \in (W, S)$. Then starting with $\tilde{F} = F_W^B$, nature could increase $\tilde{F}$ around p and correspondingly decrease it around p'. The perturbed distribution $\tilde{F}$ still satisfies the feasibility constraints, and the profit Π is reduced.

[14]If this constant were $c > 0$, then on the interval $[S, B]$ nature seeks to minimize $c \cdot \int_S^B (1 - \tilde{F}(v))\, dv$ subject to the integral inequality constraint and equal means: $\int_S^1 (1 - \tilde{F}(v))\, dv = \int_S^1 (1 - F(v))\, dv$. Thus nature equivalently maximizes $\int_B^1 (1 - \tilde{F}(v))\, dv$. Choosing $\tilde{F} = F_W^B$ results in 0 and is sub-optimal.

Hence, $r(p)$ must be supported on $[W, S]$ and $p \cdot r(p)$ is a constant. This condition together with $\int_W^S r(p)\, dp = 1$ uniquely pins down $r(p)$, which must be the density function associated with D. ■

Comparison Between Π^* and Π_{RSD}

In what follows we focus on the alternative model described in Section 1.7, where nature cannot condition on the current period price. We show that the relevant profit benchmark Π_{RSD} is in general higher than Π^*, and the difference may be significant:

Proposition 11. $\Pi_{RSD} \geq \Pi^*$ *with equality if and only if* $W = \underline{v}$ $(= p^*)$*, where* W *is as defined in the Roesler-Szentes information structure (1.3). Furthermore, as the distribution* F *varies, the ratio* Π_{RSD}/Π^* *is unbounded.*

Proof. The inequality $\Pi_{RSD} \geq \Pi^*$ is obvious. Next, recall that $\Pi^* \geq \underline{v}$ (seller can charge $\underline{v}$) and $W = \Pi_{RSD}$. Thus $W = \underline{v}$ implies $\Pi_{RSD} \leq \Pi^*$, and equality must hold.

Conversely suppose $\Pi_{RSD} = \Pi^*$, then $W = p^*(1 - G(p^*))$. This implies $p^* \geq W$. Consider a seller who charges price p^* against the Roesler-Szentes information structure F_W^B. By the unit elasticity of demand property, this seller's profit is either W (when $p^* < B$) or 0. We have shown in our one-period model that the seller can guarantee Π^* with a price of p^*. Thus the seller's profit must be W when he charges p^* and nature chooses the Roesler-Szentes information structure. Since $W = \Pi^*$ by assumption, the Roesler-Szentes information structure is a worst-case information structure for the price p^*. This yields $W \geq p^*$, because a worst-case information structure cannot include any signal that leads to a posterior expected value strictly less than p^*. We conclude $p^* = W = p^*(1 - G(p^*))$, from which it follows that $G(p^*) = 0$ and $p^* = \underline{v}$. Thus $W = \underline{v}$ must hold.

To study the ratio Π_{RSD}/Π^*, we restrict attention to a very simple class of distributions F: with probability λ, the buyer's true value is 1; otherwise her value is 0. The optimal price in the known-value case is $\hat{p} = 1$, and the corresponding profit is $\hat{\Pi} = \lambda$. In our main model, the maxmin optimal price p^* solves

$$p^* \in \operatorname*{argmax}_p \; p(1-G(p)) = \operatorname*{argmax}_{0\leq p\leq \lambda} \; p \cdot \frac{\lambda - p}{1-p}$$

Simple algebra gives $p^* = 1-\sqrt{1-\lambda}$, and $\Pi^* = (1-\sqrt{1-\lambda})^2$ which is roughly $\frac{\lambda^2}{4}$ for small λ.

Because the distribution F has two-point support, it is clear that nature can induce any $\tilde{F}$ supported on $[0,1]$ with mean λ as the distribution of posterior expected values. Thus the Roesler-Szentes information structure involves the smallest W such that F_W^B has mean λ for some $B \leq 1$. From (1.3), we compute that the mean of F_W^B is $W \log B - W \log W + W$. We look for the smallest W such that $\log B = \frac{\lambda}{W} + \log W - 1$ is non-positive. It follows that W is the smallest positive root of the equation

$$\frac{\lambda}{W} + \log W = 1.$$

For λ small, we have the approximation $\Pi_{RSD} = W \approx \frac{\lambda}{|\log \lambda|}$. Thus both ratios $\hat{\Pi}/\Pi_{RSD}$ and Π_{RSD}/Π^* are unbounded.[15] ■

Example from Section 1.9.2. Suppose $v \sim \tilde{G}$, where $\tilde{G}$ is some distribution that has expectation μ and support with upper bound $\overline{v}$. First, via similar reasoning, it is without loss to associate each signal with an action recommendation. Second, if a buyer buys, their value should be as high as possible–in this case, equal to $\overline{v}$ (which contrasts with to the setting with information). When the price is p, a buyer that does not buy should have expected value exactly equal to p. Assuming $p < \mu$ (otherwise, the seller would obtain 0 profits via null information), we have the worst case purchasing probability q therefore satisfies $q\overline{v} + (1-q)p = \mu$. Hence $q = \frac{\mu - p}{\overline{v} - p}$, and profit is $p\frac{\mu-p}{\overline{v}-p}$. The first order condition for p gives:

$$(\overline{v}-p)(\mu - 2p) + p(\mu - p) = 0 \Rightarrow p = \overline{v} - \sqrt{\overline{v}(\overline{v}-\mu)},$$

which yields profit $2\overline{v} - \mu - 2\sqrt{\overline{v}(\overline{v}-\mu)}$. Setting $\overline{v} = 1$ and $\mu = 1/2$ yields the results claimed in the main text. ■

[15]We conjecture that these profit ratios become bounded under certain regularity conditions on F.

Additional Details for all Other Sections

Proof of Claim 6. Suppose that $\mathbb{P}[p_2 < \infty] > 0$. Let the first period distribution be the same as the one that arises in the construction of Proposition 2 (recalling that this allows for random strategies by the seller), which involves $v_1 \in \{w_1, \tilde{w}\}$, which arises from a partitional information structure where w_1 is indifferent between purchasing and waiting with no further information. Consider the following choice of nature for a distribution D_{x,v_1} over values for v_2:

- If $p_2 > \tilde{w}$, then $v_2 = v_1$.
- If $p_2 \leq \tilde{w}$, then $v_2 = x^2$ with probability $\frac{v_1+x}{x+x^2}$ and $v_2 = -x$ otherwise.

Note that the buyer that arrives in the first period has a value that follows a martingale. For the second buyer, reveal no information if $p_2 \geq \tilde{w}$, and otherwise utilize the worst case partitional information structure given p_2.

The claim will follow by showing that whenever $\mathbb{P}[p_2 < \tilde{w}] > 0$, the seller's profit from the first buyer can be made arbitrarily small, since the profit from the second buyer is at most $\delta\Pi^*$ (noting that when $p_2 \geq \tilde{w}$ with probability 1, the best the seller can do is Π^* under the constructed value process). Indeed, a buyer with value $\tilde{w}$ has expected value $\tilde{w} - p_1$ from immediate purchase, but value $\delta \cdot \frac{\tilde{w}+x}{x+x^2}(x^2 - \mathbb{E}[p_2 \mid p_1])\mathbb{P}[p_2 < \tilde{w}]$, which approaches ∞ as x approaches ∞. In contrast, since the seller charges positive prices, the probability of sale is $\frac{v_1+x}{x+x^2}$, which approaches 0 as x approaches ∞. The same holds for the buyer who has expected value w_1. Hence nature can make the probability of sale to any first period buyer arbitrarily small, thereby proving the claim. ■

A.4.2 Miscellaneous Results

Uncertainty Leads to Lower Price

We prove here that uncertainty over the information structure leads the seller to choose a lower price than if the buyer knew her value.

Proposition 12. *For any continuous distribution F, let $\hat{p}$ be an optimal monopoly price under known values:*

$$\hat{p} \in \underset{p}{\text{argmax}}\ p(1 - F(p)). \tag{A.27}$$

Then any maxmin optimal price p^* satisfies $p^* \leq \hat{p}$. Equality holds only if $p^* = \hat{p} = \underline{v}$.

Proof of Proposition 12. It suffices to show that the function $p(1 - G(p))$ strictly decreases when $p > \hat{p}$, until it reaches zero. By taking derivatives, we need to show $G(p) + pG'(p) > 1$ for $p > \hat{p}$ and $G(p) < 1$.

From definition, the lowest $G(p)$-percentile of the distribution F has expected value p. That is,

$$pG(p) = \int_0^{F^{-1}(G(p))} v dF(v), \forall p \in [\underline{v}, \mathbb{E}[v]]. \tag{A.28}$$

Differentiating both sides with respect to p, we obtain

$$G(p) + pG'(p) = \frac{\partial}{\partial p}(F^{-1}(G(p))) \cdot F^{-1}(G(p)) \cdot F'(F^{-1}(G(p))) = G'(p) \cdot F^{-1}(G(p)). \tag{A.29}$$

This enables us to write $G'(p)$ in terms of $G(p)$ as follows:

$$G'(p) = \frac{G(p)}{F^{-1}(G(p)) - p}. \tag{A.30}$$

Thus,

$$G(p) + pG'(p) = \frac{G(p) \cdot F^{-1}(G(p))}{F^{-1}(G(p)) - p}. \tag{A.31}$$

We need to show that the RHS above is greater than 1, or that $F^{-1}(G(p)) < \frac{p}{1-G(p)}$ whenever $p > \hat{p}$ and $G(p) < 1$. This is equivalent to $G(p) < F(\frac{p}{1-G(p)})$, which in turn is equivalent to

$$\frac{p}{1 - G(p)} \cdot \left(1 - F\left(\frac{p}{1 - G(p)}\right)\right) < p. \tag{A.32}$$

From the definition of $\hat{p}$, we see that the LHS above is at most $\hat{p}(1 - F(\hat{p})) \leq \hat{p} < p$, as we claim to show. Moreover, when $\hat{p} > \underline{v}$, the last inequality $\hat{p}(1 - F(\hat{p})) < \hat{p}$ is strict. Tracing back the previous arguments, we see that $G(p) + pG'(p) > 1$ holds even at $p = \hat{p}$. In that case we would have the strict inequality $p^* < \hat{p}$ as desired. ■

Alternative Interpretation

In this section, we consider an information acquisition game for which our solution is of interest. The motivation borrows heavily from Roesler and Szentes (2017). They consider a game with the following timing:

- The buyer first chooses an information structure $\mathcal{I} : V \to \Delta(S)$.
- The seller then chooses a price $p \in \mathbb{R}$.
- The buyer finally decides whether or not to purchase the object.

It turns out that the resulting information structure results in a payoff for the seller that is optimal given that the information structure is the worst possible, *and assuming that the information structure does not depend on the price*.

For our setting, first consider the $T = 1$ case, and modify the Roesler-Szentes (2017) game so that the buyer's inforamtion structure can depend on the price. That is, we take the same timing as above, but allow for information structures of the form $\mathcal{I}(p) : V \to \Delta(S)$. While in practice it may be difficult to assume that the seller literally chooses this information structure, the information may indeed be provided by a third party (such as Amazon) who would have this power and potentially this objective as well.

Recall Π^* is the seller's maxmin payoff.

Proposition 13. *Consider a one-period setting where the buyer (or third party who acts to maximize the buyer's payoff) chooses the information structure. The seller's profit in the equilibrium of this game is equal to* Π^*.

Note that this does not say the equilibrium is payoff equivalent for the *buyer*. In general it will not be. Still, if one were interested in buyer payoffs, the proposition shows that our analysis is relevant for this case as well. Furthermore, the proof is relatively straightforward.

Proof of Proposition 13. We demonstrate that the optimal choice of information for the buyer results in payoff of Π^* for the seller.

Denote by $\mathcal{I}_\emptyset$ the completely uninformative information structure, and let $\mathcal{I}^*(p)$ be the worst case information structure for the seller when the price chosen is p. The buyer chooses the following information structure:

$$\mathcal{I}(p) = \begin{cases} \mathcal{I}_\emptyset, & p = \Pi^* \\ \mathcal{I}^*(p) & p \neq \Pi^* \end{cases}.$$

The seller chooses price Π^*.

First we compute the profits of the buyer and the seller. Since $\Pi^* < \mathbb{E}_{v \sim F}[v]$ whenever F is non-degenerate, trade happens with probability 1. Since trade is always efficient, total surplus is $\mathbb{E}_{v \sim F}[v]$. We thus have the buyer's surplus is $\mathbb{E}_{v \sim F}[v] - \Pi^*$ and the seller's surplus is Π^*.

Suppose there were a choice of information structure for the buyer which obtained some payoff $u > \mathbb{E}_{v \sim F}[v] - \Pi^*$. Note that the seller's payoff must be at least Π^* in any equilibrium, since this is defined to be the maxmin profit. Hence we have that total surplus in this mechanism is larger than $\mathbb{E}_{v \sim F}[v]$, which is a contradiction. Therefore, the conjectured information structure is optimal.

But if the buyer obtains $\mathbb{E}_{v \sim F}[v] - \Pi^*$, the seller must obtain exactly Π^*; again, if the seller obtained more, total surplus would be larger than $\mathbb{E}_{v \sim F}[v]$, a contradiction. Hence the seller's payoff in this game is Π^*. ■

Note that, while the buyer obtains different payoffs than in the maxmin benchmark we previously analyzed, they can still be computed simply once we have found Π^*; since trade always occurs, the buyer obtains $\mathbb{E}_{v \sim F}[v] - \Pi^*$. Note that the same argument also goes through when there are production costs, provided that trade is ex-ante efficient.[16]

The same construction works for an arbitrary horizon, replacing the choice of a single price with the choice of a constant price path of Π^*. The proof is identical, considering the discounted buyer surplus instead of an individual buyer's surplus.

[16]The same equivalence fails for Roesler and Szentes (2017) in general when there are costs.

Buyer Uncertainty

One may wonder why the sellers in our model are so much better informed than the buyers, particularly over information in the far future. While this concern may also apply to much of the literature on robustness in mechanism design, we admit that the reader may find it particularly salient here since buyer knowledge of all future information buyers is particularly demanding.

Part of the difficulty in allowing for buyer learning in contrast to seller learning is that non-Bayesian updating is significantly more complicated (which is also why the commitment benchmark is a more natural starting point than non-commitment). Nevertheless, we seek to accommodate this as follows. Recall that a *dynamic information structure* $\mathcal{I}$ is a sequence of signal sets S_t and probability distributions $I_{a,t} : R_+ \times S_a^{t-1} \times P^t \to \Delta(S_t)$, for $1 \leq t \leq T$. We now assume that nature has the ability to choose a set Ω, which we take to be a finite set, and dynamic information structures $\{\mathcal{I}^\omega\}_{\omega \in \Omega}$, with associated probability distributions $I_t^\omega : R_+ \times S^{t-1} \times P^t \to \Delta(S_t)$. Crucially, note that the signal set *does not* depend on ω, as the interpretation is that the buyer observes signal s_t in period t. Let us denote by $\Omega(s^t)$ the set of information structures the buyer believes is feasible after signal history s_t. This allows us to describe which information structures the buyer can "rule out" over time.

For simplicity, we assume that the buyer utilizes a maxmin objective, though similar arguments would apply for general uncertainty averse preferences. It is easy to show that without any restrictions, allowing for arbitrary uncertainty aversion can eliminate any seller profits.

Proposition 14. *Suppose nature can choose any set of dynamic information structures. Then the seller's worst-case single period profit (as well as dynamic profit) is equal to 0.*

Proof. Consider any price path $(p_t)_{t=1}^T$, and let $\underline{p} = \min_t p_t$. Suppose $\underline{p} > 0$. Let nature choose $S_1 = \{0,1\}, \Omega = \{0,1\}$.

- At time 1, if $\omega = 1$, the buyer observes $s_1 = 1$ if $v > \underline{p}$ and $s_1 = 0$ otherwise.
- At time 1, if $\omega = 0$, the buyer observes $s_1 = 0$ if $v > \underline{p}$ and $s_1 = 1$ otherwise.

In this case, since the buyer is maxmin, the expected payoff from purchasing is negative, hence they will never purchase. If $\underline{p} = 0$, define $\tilde{p} = \min_{p>0} p_t$. Repeating the same argument above, replacing $\underline{p}$ with $\tilde{p}$, completes the proof. ■

This result is straightforward, particularly since the martingale condition of expectations is what allowed us to avoid degenerate solutions in the first place. Hence some restriction must be made an updating, and the following appears to be the most-reasonable.

Definition 3. *The buyer is a* ***within-period Bayesian*** *if all information structures* $I_t^{\omega}(v, s^{t-1}, p^t)$ *with* $\omega \in \Omega(s^{t-1})$ *are identical.*

In words, a within-period Bayesian does not face any uncertainty over the information they have obtained up until that time. However, they may be non-Bayesian over information they receive in the future. This rules out creating arbitrarily pessimistic Hence information still has the same bite in terms of avoiding degeneracy, but no longer imposing that the buyer has significantly more knowledge of future information than the seller.

Under the within-period Bayesian assumption, Proposition 2 carries through identically, since this proposition only requires us to construct an information structure that holds the seller down to the one-period benchmark. We can also show the following:

Proposition 15. *Under the assumption that all buyers are within period Bayesians, the optimal seller strategy is a constant price path of* p^**, delivering discounted profits of* $\Pi^* \frac{1-\delta^T}{1-\delta}$.

Proof of Proposition 15. We consider the same replacement as in Proposition 2, except we assume in addition that nature reveals the true information structure to a buyer at time 1 (in addition to pushing signals forward). Obedience and algebra still demonstrate that:

$$\mathbb{E}[v \mid s_1] \leq \mathbb{E}^{s_2,\cdots,s_T}\left[\delta^{\tau-1}\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau] + (1-\delta^{\tau-1})p_1\right].$$

Since $\mathbb{E}[v \mid s_1] = \mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau]\right]$ in the case where there is no uncertainty over future information, we have $\mathbb{E}[v \mid s_1] \geq \mathbb{E}^{s_2,\cdots,s_T}\left[\mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau]\right]$. Hence we still obtain the crucial inequality:

$$p_1 \geq \frac{\mathbb{E}^{s_2,\cdots,s_T}[(1-\delta^{\tau-1}) \cdot \mathbb{E}[v \mid s_1, s_2, \cdots, s_\tau]]}{\mathbb{E}^{s_2,\cdots,s_T}[1-\delta^{\tau-1}]}.$$

Since the buyer is a Bayesian in the transformed setting, the same analysis applies and we have $p_1 \geq \mathbb{E}[v \mid s_1, \underline{s}]$. Hence the replacement is obedient, so the probability sale is $\mathbb{E}[\delta^{t-1}]$, according to the realized information structure (noting that this is the parameter that matters for the seller's profit). Hence the seller is made no better off with the replacement. ■

We make two observations on general features of this model, where we do not restrict buyer uncertainty to being of the maxmin variety: First, our result that the worst-case for the seller is a Bayesian buyer under within-period Bayesianism should hold under more general models of ambiguity aversion than the one we present here. Indeed, Riedel (2009) derives a version of the Optional Sampling Theorem for multiple-prior supermartingales, in dynamic models of ambiguity where the agent is time-consistent.

Second, we should *not* expect our results to hold in cases where the buyer is ambiguity *loving* but a within-period Bayesian. In that case, even in the case where the buyer is a within-period Bayesian, nature can induce delay by utilizing ambiguity over the buyer's future information. In that case, preventing the buyer from purchasing in the future by raising the price may help profits, even potentially at the expense of excluding future buyers, in contrast to our main results.

Appendix B

Appendix to Chapter 2

B.1 Appendix

B.1.1 Omitted Proofs

Proof of Proposition 4. Since the developer's cost function is convex, the developer's payoff function is a Fenchel transformation of a convex function and is hence convex itself. It follows that $\hat{p}(y)\pi_R(1) - \pi_R(\hat{p})$ is a measure of uncertainty according to Ely, Frankel and Kamenica (2015), since it is 0 at degenrate beliefs and concave for all interior beliefs (since it is a linear function minus a convex function), as desired. □

Proof of Lemma 4. We find it useful to demonstrate the following, more general equation, for arbitrary function $t(p)$ and $f(p)$, noting that Lemma 4 follows immediately by setting $t(p) = \lambda e(p) + g(p)$ and $f(p) = g(p)$: For functions $t, f : [0,1] \to \mathbb{R}$,

$$p_0 \mathbb{E}_y[t(\hat{p}(y)) \mid T] + (1 - p_0)\mathbb{E}_y[f(\hat{p}(y)) \mid F] = \mathbb{E}_y[\hat{p}(y)t(\hat{p}(y)) + (1 - \hat{p}(y))f(\hat{p}(y))] \quad \text{(B.1)}$$

By writing out the definition of the conditional expectation, we have:

$$\begin{aligned} p_0 \mathbb{E}_y[t(\hat{p}(y)) \mid T] &= p_0 \sum_{y \in Y} t(\hat{p}(y)) \mathbb{P}[y \mid T] \\ &= \sum_{y \in Y} t(\hat{p}(y)) \left(p_0 \frac{\mathbb{P}[y \mid T]}{\mathbb{P}[y]} \right) \mathbb{P}[y] \\ &= \sum_{y \in Y} t(\hat{p}(y)) \hat{p}(y) \mathbb{P}[y] = \mathbb{E}_y[\hat{p}(y) t(\hat{p}(y))]. \end{aligned}$$

An almost identical argument can be used for the other term in (B.1). In fact, this generalizes for any number of states, not just $\theta \in \{T, F\}$. □

Proof of Lemma 5. **Proof of (1)** Suppose M is the index set of observable indices, and partition the scientist's action into $a = (a_M, a_{-M})$. We show that there is some a^*_{-M} such that when the developer conjectures that a^*_{-M} are the unobserved actions of the scientist, the scientist's best response is to follow action a^*_{-M}. Since p_0 is interior and $h_T(a) > h_F(a)$ for all a, the developer always puts non-negative probability on observing $y = 0$ or $y = 1$, for any conjecture regarding the scientist's behavior. Therefore, there are unique beliefs $(\hat{p}(0), \hat{p}(1))$ formed after observing a signal $y = 0$ or $y = 1$, respectively, for any equilibrium strategy of the scientist. In fact, given a conjecture, since A is compact, we have that $\mathbb{P}[y = 0]$ and $\mathbb{P}[y = 1]$ are both bounded away from 0. This implies that beliefs are a continuous function of actions, and well-defined given any conjecture.

Let $t(p) = \lambda e(p) + g(p)$ and $f(p) = g(p)$, and define the function ϕ as follows:

$$\phi(a_{-M}) = \arg \max_{\tilde{a}_{-M} \in A_{-M}} p_0[t(\hat{p}_{(a_M, a_{-M})}(0)) + h_T(a_M, \tilde{a}_{-M})(t(\hat{p}_{(a_M, a_{-M})}(1)) - t(\hat{p}_{(a_M, a_{-M})}(0)))]$$
$$+ (1 - p_0)[f(\hat{p}_{(a_M, a_{-M})}(0)) + h_F(a_M, \tilde{a}_{-M})(f(\hat{p}_{(a_M, a_{-M})}(1)) - f(\hat{p}_{(a_M, a_{-M})}(0)))] - c(a_M, a_{-M})$$

Note that $\phi(a_{-M})$ gives the payoff maximizing response, assuming (observable) actions a_M are chosen and a conjecture of a_{-M}. Taking $a^n_{-M} \to a_{-M}$, and $b_n \in \phi(a_{-M})$ with $b_n \to b$, since beliefs are continuous in a and $f(p), t(p)$ are continuous as well (by continuity of $e(p)$), we have

$$t(\hat{p}_{(a_M, a^n_{-M})}(1)) - t(\hat{p}_{(a_M, a^n_{-M})}(0)) \to t(\hat{p}_{(a_M, a_{-M})}(1)) - t(\hat{p}_{(a_M, a_{-M})}(0)),$$

and similarly for f. If $b \notin \phi(a_{-M})$, then there exists some value δ such that:

$$p_0(h_T(a_M,\delta) - h_T(a_M,b))(t(\hat{p}_{(a_M,a_{-M})}(1)) - t(\hat{p}_{(a_M,a_{-M})}(0)))$$
$$+ (1-p_0)(h_F(a_M,\delta) - h_F(a_M,b))(f(\hat{p}_{(a_M,a_{-M})}(1)) - f(\hat{p}_{(a_M,a_{-M})}(0))) > c(a_M,\delta) - c(a_M,b).$$

But since $a^n_{-M} \to a_{-M}$ and $b_n \to b$, by continuity we can find some n sufficiently large such that

$$p_0(h_T(a_M,\delta) - h_T(a_M,b_n))(t(\hat{p}_{(a_M,a^n_{-M})}(1)) - t(\hat{p}_{(a_M,a^n_{-M})}(0)))$$
$$+ (1-p_0)(h_F(a_M,\delta) - h_F(a_M,b_n))(f(\hat{p}_{(a_M,a^n_{-M})}(1)) - f(\hat{p}_{(a_M,a^n_{-M})}(0))) > c(a_M,\delta) - c(a_M,b_N),$$

contradicting that b_n is a maximizer of $\phi(a^n_{-M})$. Hence the map ϕ is upper-hemicontinuous.

Furthermore, $\phi(a_{-M})$ is nonempty and closed because A_{-M} is compact (being the product of intervals) and the objective function in the expression for $\phi(a_{-M})$ is continuous. Finally, to see that it is convex, notice that if a'_{-M}, a''_{-M} are both in $\phi(a_{-M})$ the convexity of $c_S(a)$ and the concavity of $h_T(a), h_F(a)$ means that we must have that

$$p_0 h_T(\tilde{a}_{-M})(t(\hat{p}_{(a_M,a_{-M})}(1)) - t(\hat{p}_{(\alpha,a_{-M})}(0)))$$
$$+ (1-p_0)h_F(a_M,\tilde{a}_{-M})(f(\hat{p}_{(a_M,a_{-M})}(1)) - f(\hat{p}_{(a_M,a_{-M})}(0))) - c_\alpha(\tilde{a}_{-M}).$$

is constant for all $\tilde{a}_{-M}$ with $\tilde{a}_{-M} = \alpha a'_{-M} + (1-\alpha)a''_{-M}$ with $\alpha \in [0,1]$, so that $\phi(a_{-M})$ is convex. Hence by Kakutani's fixed point theorem, an equilibrium exists when a_M is observed, for any choice of a_M.

That the first order condition is given by Equation (2.3) follows from observing that $\hat{p}_{(a_M,a_{-M})}$ does not respond to the choice of a_i for $i \notin M$. In that case, we see that the marginal benefit is given by:

$$p_0 \frac{\partial h_T(a_M,a_{-M})}{\partial a_i}(t(\hat{p}_{(a_M,a^n_{-M})}(1)) - t(\hat{p}_{(a_M,a^n_{-M})}(0))) + (1-p_0)\frac{\partial h_F(a_M,a_{-M})}{\partial a_i} f(\hat{p}_{(a_M,a^n_{-M})}(1)) - f(\hat{p}_{(a_M,a^n_{-M})}(0)),$$

which reduces to the stated condition when equal to marginal cost, as desired.

Proof of (2) Suppose to the contrary that there is a mixed strategy equilibrium. Then the first order condition in equation (2.3) must hold for two values of a_{-M}, say $a^1_{-M} < a^2_{-M}$. On the other hand, the developer's beliefs do not depend on the choice of a_{-M}. Keeping the notation from the previous we have:

$$\nabla_{a_{-M}} c(a_M, a^i_{-M}) = p_0(t(\hat{p}(1)) - t(\hat{p}(0)))\nabla_{a_{-M}} h_T(a_M, a^i_{-M}) \\ + (1-p_0)(f(\hat{p}(1)) - f(\hat{p}(0)))\nabla_{a_{-M}} h_F(a_M, a^i_{-M}),$$

and hence subtracting the equation for $i = 1$ from the equation for $i = 2$, and taking the dot product for some arbitrary α with $||\alpha|| = 1$, we have:

$$\alpha \cdot \nabla_{a_{-M}}(c(a_M, a^2_{-M}) - c(a_M, a^1_{-M})) = p_0(t(\hat{p}(1)) - t(\hat{p}(0)))\alpha \cdot \nabla_{a_{-M}}(h_T(a_M, a^2_{-M}) - h_T(a_M, a^1_{-M})) \\ + (1-p_0)(f(\hat{p}(1)) - f(\hat{p}(0)))\alpha \cdot \nabla_{a_{-M}}(h_F(a_M, a^2_{-M}) - h_F(a_M, a^1_{-M})).$$

By the multivariate mean value theorem, applied to h_T, h_F and c, we have, for some a_t, a_f and a_c which are convex combinations of a^1_{-M} and a^2_{-M} such that:

$$\alpha \cdot \nabla^2_{a_{-M}} c(a_M, a_c)(a^2_{-M} - a^1_{-M}) = p_0(t(\hat{p}(1)) - t(\hat{p}(0)))\alpha \cdot \nabla^2_{a_{-M}} h_T(a_M, a_t)(a^2_{-M} - a^1_{-M}) \\ + (1-p_0)(f(\hat{p}(1)) - f(\hat{p}(0)))\alpha \cdot \nabla^2_{a_{-M}} h_F(a_M, a_f)(a^2_{-M} - a^1_{-M}).$$

But since either h_T or h_F is strictly concave or c is strictly convex, either the left hand side is strictly positive or the right hand side is strictly negative, with both being at least weakly so, a contradiction. Hence in equilibrium, there can only be pure strategies. □

Proof of Corollary 2. This follows from demosntrating that the scientist's payoffs, as a function of the developer's beliefs, are convex. This is immediate in the case of polynomial effort costs. Indeed, we take the second derivative of $pe(p)$ (the benefit from follow-on effort) and observe that it is equal to:

$$\lambda(2e'(p) + e''(p)).$$

Since $c_R(e)$ is strictly convex, $e'(p) > 0$, since the first order condition is:

$$bp = c'_R(e(p)),$$

and differentiating with respect to p gives:

$$b = c_R''(e(p))e'(p),$$

and differentiating again gives:

$$0 = c_R'''(e)(e'(p))^2 + c_R''(e)e''(p).$$

Since $e(p)$ is strictly increasing, the assumptions on $c_R'''(e)$ ensure that $e''(p) \geq 0$, and hence the objective is convex. □

In general, convexity of developer effort by itself is not a strong enough assumption in order to ensure that $pe(p)$ is convex. To see this, suppose that:

$$c_R(e) = 1 - \sqrt{1-e} \Rightarrow c_R'(e) = \frac{1}{2\sqrt{1-e}} > 0 \Rightarrow c_R''(e) = \frac{1}{4(1-e)^{3/2}} > 0.$$

In that case, we have:

$$e(p) = \max\{0, -\frac{1}{4p^2} + 1\},$$

and we observe that $pe(p)$ is concave whenever $e(p) > 0$.

Proof of Theorem 7. Consider the following family of cost functions:

$$c_{\eta,\gamma,\tilde{a}}(a) = \eta \sum_{j \in J}(a_j - \underline{a}_j)^2 + \gamma \sum_{k \in K}(a_k - \underline{a}_k)^2(a_j - \tilde{a}_j)\mathbf{1}[a_j \geq \tilde{a}_j],$$

for $\eta, \gamma > 0$ and $\tilde{a}_j > \underline{a}_j$. Note that these are strictly convex in a_K for all parameters, since it is a quadratic function when we treat a_j as a constant. By Corollary 2, when distortive action a_k is observed, it is set equal to $\underline{a}_k$ in equilibrium. Let η^* denote the smallest value of η such that $a_j = \underline{a}_j$ for all informative actions. Since all distortive actions are set equal to $\underline{a}_k$ under complete observability, η^* in does not depend on the choice of γ or $\tilde{a}$.

Now suppose that the K coordinates are unobserved. Note that for all $k \in K$, we have:

$$\frac{\partial c_{\eta^*,\gamma}}{\partial a_k} = 2\gamma(a_k - \underline{a}_k) + \sum_{j \in J}(a_j - \underline{a}_j)\mathbf{1}[a_j \geq \tilde{a}_j].$$

Note that at $a_i = \underline{a}_i$ for all actions i (as is the scientist's choice under observability, by construction), the first order condition in (2.3) cannot be satisfied for any $\gamma > 0$. It follows that some distortive dimension a_k must be in the interior (or possibly at $\overline{a}_k$) whenever γ is positive.

Now let $\gamma \to \infty$ and set $\tilde{a}_j = \underline{a}_j + \varepsilon$. We claim that for all j, $a_j \to a^* \geq \tilde{a}_j > \underline{a}_j$ (passing to a subsequence if necessary, we can ensured that the limit always exists since the action space is compact) when ε is sufficiently small. Suppose otherwise; in that case, $a_k \to \overline{a}^k$ by Lemma 5, since the marginal cost is 0. It follows that the loss to the scientist due to distortion does not vanish as $\gamma \to \infty$, for every $a_k \in [\underline{a}_k, \underline{a}_k + \varepsilon]$ and for any ε. On the other hand, by taking ε sufficiently small, we can ensure that the cost of raising the action for $\underline{a}_j$ to $\tilde{a}_j$ is negligible, in which case the scientist would find it to be an equilibrium to choose $a_k = \underline{a}_k$. It follows that there exists a pair (γ, ε) such that making the K coordinates unobserved results in a more informative experiment chosen than under full transparency. □

Proof of Proposition 6. Without costs, then under the restrictions of preferences studied in this paper, the scientist's payoff is increasing in the informativeness of the experiment. Hence the scientist chooses the maximally informative experiment under observability. On the other hand, making any biasing dimension unobservable results in $a_i = \overline{a}_i$ being chosen by the scientist in equilibrium, by Lemma 5. The result follows from observing that by definition, the most informative experiment under $a_{-M} = \underline{a}_{-M}$ is less informative thatn the most informative experiment under $a_{-M} = \overline{a}_{-M}$. □

B.1.2 Scientist Preferences over y

In this appendix, we comment on a modification to the model where we allow for the scientist to have preferences over y itself. While this could take several different forms in general, for simplicity we will comment on the case where the payoffs are separable, and the scientist obtains an added benefit of $\lambda_y \cdot y$. from a positive result.

In general this does not interfere with our application of the belief-based approach, noting that any positive result leads to a higher belief and any negative result leads to a lower belief. Hence this setting is as if there were a jump in the scientist's payoff function at the prior. That said, it is simplest to comment on this case simply by inspection. In this case, it is immediate that the scientist is incentivized to maximize the biasing action in this case (whether higher informative actions will be taken depends on the prior):

Proposition 16. *There exists $\overline{\lambda}_y$ such that when $\lambda_y > \overline{\lambda}_y$, biasing actions are all maximized in equilibrium.*

More to the point of the paper, however, by itself transparency does not interact with the experiment choice when these kinds of considerations are dominant:

Proposition 17. *Suppose $\lambda = g(p) = 0$ and $\lambda_y > 0$. Then the scientist's experiment choice does not differ depending on transparency or not.*

This is immediate since the prior does not depend on experiment choices in this setting.

While immediate, the result may be counterintuitive given that it is natural to feel inclined to prefer positive results rather than negative results. One may be that negative results may be harder to publish than positive results, in which case there should actually be an *interaction* between positive results and the other payoff terms. One could accomodate this and obtain similar results, as this is similar to imposing even greater convexity in the payoffs. Alternatively, it may be that positive results that are obtained "cheaply" (via bias) are less meaningful, but those that are achieved "scrupulously" (via informativeness) are more meaningful. This would suggest greater interdependence between the cost function and the benefit than what we have here. Since this would take us too far afield, we do not pursue such specifications further.

Appendix C

Appendix to Chapter 3

C.1 Appendix

Proof of Lemma 6. As stated in the main text, we replace a mechanism that possibly correlates the project choice with the selected agent with one that "simulates out" the impact of other agents on prototype choice. We illustrate the argument for the case where there is a single risky prototype; the case with an arbitrary number of prototypes is analogous, but the notation is more involved. With this simplification, we can write an allocation rule by $\{Y_i^S(\theta_i, \theta_{-i}), Y_i^R(\theta_i, \theta_{-i})\}$, where $Y_i^j(\theta_i, \theta_{-i})$ is the probability agent i is asked to develop a particular project j, where $j = S$ denotes safe and $j = R$ denotes risky, given the vector of reports (θ_i, θ_{-i}).

Note that the agent's payoff from announcing $\hat{\theta}_i$ when they are of type θ_i can be written as

$$\mathbb{E}_{\theta_{-i}}[Y_i^S(\hat{\theta}_i, \theta_{-i}) v_S^A(0) + Y_i^R(\hat{\theta}_i, \theta_{-i})(\theta_i v_S^A(x_i) + (1 - \theta_i) v_F^A(x_i))].$$

Define:

$$Y_i(\hat{\theta}_i, \theta_{-i}) = Y_i^S(\hat{\theta}_i, \theta_{-i}) + Y_i^R(\hat{\theta}_i, \theta_{-i}), \text{ and } X^j(\hat{\theta}_i) = \frac{\mathbb{E}_{\theta_{-i}}[Y_i^j(\hat{\theta}_i, \theta_{-i})]}{\mathbb{E}_{\theta_{-i}}[Y_i(\hat{\theta}_i, \theta_{-i})]}.$$

Now consider the mechanism that awards the contract to agent i with probability $Y_i(\hat{\theta}_i, \theta_{-i})$, and develops project j with probability $X_i^j(\hat{\theta}_i)$. Since $\mathbb{E}_{\theta_{-i}}[Y_i^j(\hat{\theta}_i, \theta_{-i})] = \mathbb{E}_{\theta_{-i}}[Y_i(\hat{\theta}_i, \theta_{-i})] X^j(\hat{\theta}_i)$, this mechanism gives the agent exactly the same interim utility as the original mechanism. Noting that (now assuming incentive compatibility) we can write the principal's payoff as:

$$\sum_{i=1}^{N} \mathbb{E}_{\theta_i} \left[\mathbb{E}_{\theta_{-i}} \left[Y_i^S(\theta_i, \theta_{-i}) v_S^P(0) + Y_i^R(\theta_i, \theta_{-i})(\theta_i v_S^P(x_i) + (1-\theta_i) v_F^P(x_i)) \right] \right],$$

we see that each term inside the expectation $\mathbb{E}_{\theta_i}$ is the same in the new mechanism, and hence we have not changed the principal's payoffs, either. It follows that letting the allocation depend only on the selected agent's private information is without loss. □

We can also show the following Lemma of independent interest. The argument is similar to the proof of the taxation principle in the case where transfers are present:

Lemma 15. *Let $\theta^i, \tilde{\theta}^i$ differ in a single coordinate j, with $\theta^i_j > \tilde{\theta}^i_j$. Then in any incentive compatible mechanism, $\mathbb{E}_{\theta^{-i}}[Y_i(\theta^i, \theta^{-i})] X(\theta^i)[x_j] \geq \mathbb{E}_{\theta^{-i}}[Y_i(\tilde{\theta}^i, \theta^{-i})] X(\tilde{\theta}^i)[x_j]$.*

In words, it follows from incentive compatibility that both the probability of winning and the probability of being allocated a given allocation is weakly increasing in each prototype quality type.

Proof. Note that, since type θ does not mimic type $\tilde{\theta}$ in the mechanism X, Y, incentive compatibility implies:

$$\mathbb{E}_{\theta^{-i}}[Y_i(\theta^i, \theta^{-i})] \sum_{x_j} X(\theta^i)[x_j](\theta^i_j v^1_A(x_j) + (1-\theta^i_j) v^0_A(x_j)) \geq$$

$$\mathbb{E}_{\theta^{-i}}[Y_i(\tilde{\theta}^i, \theta^{-i})] \sum_{x_j} X(\tilde{\theta}^i)[x_j](\theta^i_j v^1_A(x_j) + (1-\theta^i_j) v^0_A(x_j)).$$

Similarly,

$$\mathbb{E}_{\theta^{-i}}[Y_i(\tilde{\theta}^i, \theta^{-i})] \sum_{x_j} X(\tilde{\theta}^i)[x_j](\tilde{\theta}^i_j v^1_A(x_j) + (1-\tilde{\theta}^i_j) v^0_A(x_j)) \geq$$

$$\mathbb{E}_{\theta^{-i}}[Y_i(\theta^i, \theta^{-i})] \sum_{x_j} X(\theta^i)[x_j](\tilde{\theta}^i_j v^1_A(x_j) + (1-\tilde{\theta}^i_j) v^0_A(x_j)).$$

Now suppose that $\tilde{\theta}$ and θ are equal in all coordinates other than the jth coordinate. In this case, adding the inequalities together yields:

$$\mathbb{E}_{\theta^{-i}}[Y_i(\theta^i,\theta^{-i})]X(\theta^i)[x_j]((\theta^i_j-\tilde{\theta}^i_j)(v^1_A(x_j)-v^0_A(x_j))\geq \mathbb{E}_{\theta^{-i}}[Y_i(\tilde{\theta}^i,\theta^{-i})]X(\tilde{\theta}^i)[x_j]((\theta^i_j-\tilde{\theta}^i_j)(v^1_A(x_j)-v^0_A(x_j)).$$

Hence if $\theta^i_j > \tilde{\theta}^i_j$, it follows that:

$$\mathbb{E}_{\theta^{-i}}[Y_i(\theta^i,\theta^{-i})]X(\theta^i)[x_j]\geq \mathbb{E}_{\theta^{-i}}[Y_i(\tilde{\theta}^i,\theta^{-i})]X(\tilde{\theta}^i)[x_j],$$

as desired. □

For the following proof, recall the definition of the efficient prototype as the one that is optimal for both parties when the alignment assumption holds.

Proof of Lemma 7. Fix any incentive compatible allocation rule. We can write incentive compatibility as:

$$\mathbb{E}[Y_i(\theta^i,\theta^{-i})]\sum_{x_j}X(\theta^i)[x_j](\theta^i_j v^1_A(x_j)+(1-\theta^i_j)v^0_A(x_j))\geq$$
$$\mathbb{E}[Y_i(\tilde{\theta}^i,\theta^{-i})]\sum_{x_j}X(\tilde{\theta}^i)[x_j](\theta^i_j v^1_A(x_j)+(1-\theta^i_j)v^0_A(x_j)). \quad \text{(C.1)}$$

Consider any type such that the safe prototype is efficient. Upon inspection of (C.1), replacing a rule where all such types implement the safe prototype is still incentive compatible– for all other types for which the risky prototype is efficient, this makes mimicking even less attractive, and for all types for which the safe prototype is efficient, this again makes the incentives even stronger.

On the other hand, consider any type θ for which a risky prototype is efficient. In this case, any optimal mechanism can restrict to randomizations between the most efficient risky prototype and the safe prototype;

□

We comment on this proof in relation to Example 5.1, where the property of deterministic project choice fails even if it holds in the single agent model. In that case, shifting probability away from the third project incentivized the low type to mimic the high type. So in the example, mimicking becomes more attractive for the low type. In contrast, when the alignment assumption holds, the low type is not incentivized to mimic the high type: This type prefers the safe prototype, and increasing the probability that the good type implements the risky prototype would only make this *less* attactive.

Proof of Lemma 8. We provide the argument for the case of two types for each agent, noting that the general case is analogous simply by considering the adjacent types. By Lemma 7, the high type will undertake the risky project and the low type will undertake the safe project. As such the incentive compatibility constraints are:

$$\mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\overline{\theta}, x_i, \theta_{-i}, x_{-i})](\overline{\theta}_{x_i} v_A^S(x_i) + (1-\overline{\theta})v_A^F(x_i)) \geq \mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i})]v_A(0)$$
$$\mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i})]v_A(0) \geq \mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\overline{\theta}, x_i, \theta_{-i}, x_{-i})](\underline{\theta}_{x_i} v_A^S(x_i) + (1-\underline{\theta}_x)v_A^F(x_i)).$$

Now, the left hand side of the second inequality and the right hand side of the first inequality are identical. On the other hand, the right hand side of the second inequality is always less than the left hand side of the first inequality. Therefore, if the first equation holds with equality, the second holds with slack, and visa versa.
Suppose the first equation held with equality. This then implies, for a positive measure (x_i, θ):

$$\mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\overline{\theta}, x_i, \theta_{-i}, x_{-i})] < \mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i})]$$
$$\Leftrightarrow \mathbb{E}_{\theta_{-i},x_{-i}}[Y_i(\overline{\theta}, x_i, \theta_{-i}, x_{-i}) - Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i})] < 0.$$

Let Z_i be the set of agent x_i types where this is strict. For each $z_i \in Z_i$, there exists a positive measure set $Z(z_i) \subset \Theta_{-i} \times X_{-i}$ where $Y_i(\overline{\theta}, x_i, \theta_{-i}, x_{-i}) < Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i})$, and $Z(z_i) = \cup_{j \neq i} Z_j(z_i)$ where j wins with positive probability in $Z_j(z_i)$. We can partition $Z_j(z_i)$ further into $Z_j^S(z_i)$ and $Z_j^R(z_i)$, where j implements "risky" on $Z_j^R(z_i)$ and "safe" on $Z_j^S(z_i)$. Suppose $Z_j^R(z_i)$ has positive probability for a positive measure subset of Z_i. Consider a new mechanism $\tilde{Y}$ which is identical for all agents not equal to i or j, and for any $z_i \in Z_i$ and

$z_j \in Z_j^R(z_i)$ sets:

$$\tilde{Y}_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) = Y_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) + \frac{\varepsilon}{\mathbb{P}[\overline{\theta}_{z_i}]}, \tilde{Y}_i(\underline{\theta}, z_i, \theta_{-i}, x_{-i}) = Y_i(\underline{\theta}, z_i, \theta_{-i}, x_{-i}) - \frac{\varepsilon}{\mathbb{P}[\underline{\theta}_{z_i}]}$$

$$\tilde{Y}_j(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) = Y_j(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) - \frac{\varepsilon}{\mathbb{P}[\overline{\theta}_{z_i}]}, \tilde{Y}_j(\underline{\theta}, z_i, \theta_{-i}, x_{-i}) = Y_j(\underline{\theta}, z_i, \theta_{-i}, x_{-i}) + \frac{\varepsilon}{\mathbb{P}[\underline{\theta}_{z_i}]}$$

First, the probability that all agents $k \neq i, j$ win are unchanged, so their incentive constraints still hold. Second, the expected utility player j obtains is proportional to their probability of winning, which is also unchanged. Third, the expected *principal* surplus for all agents *other than i* (including agent j) is unchanged as well. Fourth, this mechanism is feasible for ε sufficiently small: $Y_i(\underline{\theta}, x_i, \theta_{-i}, x_{-i}) > 0$ implies $Y_j(\underline{\theta}, z_i, \theta_{-i}, x_{-i}) < 1$, but by consturction $Y_j(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) > 0$, while we also had $Y_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) < Y_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i})$, so that $0 \leq \tilde{Y}_i \leq 1$ and likewise for $\tilde{Y}_j$. Additionally, for ε sufficiently small, agent i's incentive constraints still hold: type $\overline{\theta}$ is more likely to win, so he does better, and type $\underline{\theta}$ was slack. As for the principal, her utility from agent i under the new mechanism is

$$\mathbb{P}[\overline{\theta}_{z_i}] \left((\overline{\theta} v_S^P(x_i) + (1 - \overline{\theta}) v_F^P(x_i))(Y_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) + \frac{\varepsilon}{\mathbb{P}[\overline{\theta}_{z_i}]}) \right) + \mathbb{P}[\underline{\theta}_{z_i}] v_S^P(0) \left(Y_i(\overline{\theta}, z_i, \theta_{-i}, x_{-i}) + \frac{\varepsilon}{\mathbb{P}[\underline{\theta}_{z_i}]} \right),$$

and subtracting her payoff under the original mechanism, the change is:

$$\varepsilon \left((\overline{\theta} v_S^P(x_i) + (1 - \overline{\theta}) v_F^P(x_i)) - v_S^P(0) \right) > 0$$

and hence the principal is better off as well. This contradicts optimality, meaning that the good realization never wants to mimic the bad one, as desired. □

Proof of Proposition 9. First, note the under the assumption of symmetry, any optimal mechanism can be symmetrized so that the allocation probabilities are the same across agents. That is, we consider all possible permutation of the agents, and note that the correspondingly permuted mechanism must also be an optimal mechanism. Taking the average over all such mechanisms maintains the incentive compatibility constraints and is hence optimal as well.

Hence suppose we have some symmetric mechanism $\{Y(\theta_i, \theta_{-i}), X(\theta_i)\}$. Then for all α, multiplying $Y(\theta_i, \theta_{-i})$ by α maintains incentive compatibility.

Suppose the incentive compatibility constraints do not bind in this solution. It follows that whenever all types prefer the safe prototype, they strictly prefer to do so. Hence in this event, we can allow the $N+1$st agent to develop the risky prototype with some small probability if they would like. Since they only make this announcement when they prefer it, by the alignment assumption, this increases the all agent's payoffs.

Now suppose the incentive compatibility constraints do bind. In this case, we must have that, for any tpye θ that prefers the risky prototype, $\mathbb{E}_{\theta^{-i}}[Y(\theta, \theta^{-i})] < 1$. Hence it follows that there is some positive probability that the safe prototype is developed, even when the risky prototype is preferred by the principal. Consider the following mechanism: Award the contract according to the N agent mechanism, but whenever it would choose the safe prototype over a risky prototype, let the $N+1$st agent obtain the contract with probabilities $\alpha Y(\theta_i, \theta_{-i})$ (sticking to the original allocation if these probabilities dictate the new agent does not develop it). If the agent would obtain the contract with the risky prototype, allocate it to the other agent with probability $\alpha k Y(\theta_i, \theta_{-i})$, where k is chosen so that the incentive constraints still bind (noting that $k < 1$ by alignment).

We show that this increases the principal's profits. Letting U_P denote the principal's surplus from a single agent according to the mechanism $Y(\theta_i, \theta_{-i})$, we have the principal gains $\alpha\mathbb{P}[\text{Safe wins}](\mathbb{E}[U_P] - v_P(0))$, but loses $\alpha k \mathbb{P}[\text{Risky wins}](\theta v_P^1(x_j) + (1-\theta)v_P^0(x_j) - \mathbb{E}[U_P])$. We wish to show that:

$$\mathbb{P}[\text{Safe wins}](\mathbb{E}[U_P] - v_P(0)) > k\mathbb{P}[\text{Risky wins}](\theta v_P^1(x_j) + (1-\theta)v_P^0(x_j) - \mathbb{E}[U_P]).$$

Note, however, that using the definition of U_P, it is immediate that:

$$\mathbb{P}[\text{Safe wins}](\mathbb{E}[U_P] - v_P(0)) = \mathbb{P}[\text{Risky wins}](\theta v_P^1(x_j) + (1-\theta)v_P^0(x_j) - \mathbb{E}[U_P])$$

so that the desired inequality follows due to the observation that $k < 1$, as desired. □

Proof of Proposition 10. Suppose all agents follow a strategy that announce the risky prototype whenever doing so is efficient. Consider the agent with the most efficient risky prototype (given the expected type conditional on that prototype being better). It follows that this agent wins with probability 1 whenever proposing the risky prototype, since this delivers the highest ex-post payoffs to the principal. However, by assumption, the optimal mechanism with commitment does not allocate the contract to this agent with probability 1. It follows that this cannot be incentive compatible, since otherwise always allocating this contract in the commitment case would be efficient, as desired. □

C.2 Microfounding Alignment

In this section, we consider how the alignment assumption may arise when the agent has private information about cost after observing whether the prototype is feasible. Suppose the cost of prototype x, which delivers value $v(x)$ is $c \sim F_{v(x)}$[1] Denote by $\phi_v(c)$ the corresponding virtual value, $c + \frac{F(c)}{f(c)}$. Standard analysis shows that the interim optimal mechanism involves a price paid to procure the prototype of $\phi_v^{-1}(v)$. Henceforth, we take as given that the principal will utilize such a transfer (noting that if alignment holds, this will indeed be optimal, but if it does not then it may be optimal to overpay or underpay to align incentives).

We assume that the firm sets a price of $\phi^{-1}(v)$ when the prototype is of value v.[2] In this case, the value to prototype v to principal and agent, respectively, is:

$$\pi_P(v) = F(\phi_v^{-1}(v))(v - \phi_v^{-1}(v)),$$
$$\pi_A(v) = \int_0^{\phi_v^{-1}(v)} (\phi_v^{-1}(v) - c) f(c) dc = \phi_v^{-1}(v) F(\phi_v^{-1}(v)) - \int_0^{\phi_v^{-1}(v)} c f_v(c) dc.$$

We first differentiate the agent's payoff with respect to v:

$$\begin{aligned}\frac{d\pi_A(v)}{dv} &= \frac{d(\phi^{-1}(v))}{dv} F(\phi^{-1}(v)) + \phi^{-1}(v) f(\phi^{-1}(v)) \frac{d(\phi^{-1}(v))}{dv} - \phi_v^{-1}(v) f(\phi^{-1}(v)) \frac{d(\phi^{-1}(v))}{dv} \\ &= \frac{d(\phi^{-1}(v))}{dv} F(\phi^{-1}(v)).\end{aligned}$$

[1]It is helpful to refer to cost distributions as being indexed by values and not by projects, as we will see below.

[2]This is not necessarily obvious, as one has to think about what price can be charged in the interim/ex-ante/etc., and so in the general model it could easily be something different.

On the other hand,

$$\frac{d\pi_P(v)}{dv} = F(\phi^{-1}(v))(1 - \frac{d(\phi^{-1}(v))}{dv}) + (v - \phi^{-1}(v))f(\phi^{-1}(v))\frac{d(\phi^{-1}(v))}{dv}.$$

We factor out $\frac{d\pi_A(v)}{dv}$ from this expression:

$$\frac{d\pi_P(v)}{dv} = \frac{d\pi_A(v)}{dv}\left(\left(\frac{1}{\frac{d(\phi^{-1}(v))}{dv}} - 1\right) + (v - \phi^{-1}(v))\frac{f(\phi^{-1}(v))}{F(\phi^{-1}(v))}\right)$$

Now, notice that $\phi(x) = x + \frac{F(x)}{f(x)}$. Hence $\frac{f(x)}{F(x)} = \frac{1}{\phi(x)-x}$. We therefore make this substitution into the above formula and obtain:

$$\frac{d\pi_P(v)}{dv} = \frac{d\pi_A(v)}{dv}\left(\frac{1}{\frac{d(\phi^{-1}(v))}{dv}}\right)$$

and therefore:

$$\frac{d(\phi^{-1}(v))}{dv} = \frac{\frac{d\pi_A(v)}{dv}}{\frac{d\pi_P(v)}{dv}}$$

as desired.

Hence, when ϕ is linear, $\frac{d\phi^{-1}(v)}{dv}$ is constant, and hence the principal's payoff is linear in the agent's payoff, and hence alignment holds. If $\frac{d\phi^{-1}(v)}{dv}$ is increasing, then the agent's payoff using the (Myerson optimal) transers mentioned above increases more quickly than the principal's payoff, and hence they prefer riskier projects than the principal. The opposite conclusion holds when $\frac{d\phi^{-1}(v)}{dv}$ is decreasing.