



Principal-Agent Problems and Experimentation

Citation

Clark, Aubrey Benedict. 2017. Principal-Agent Problems and Experimentation. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:42061499>

Terms of Use

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

Share Your Story

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

[Accessibility](#)

Principal-Agent Problems and Experimentation

A dissertation presented

by

Aubrey Clark

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

June 2017

© 2017 Aubrey Clark

All rights reserved.

Dissertation Advisor:
Professor Eric Maskin

Author:
Aubrey Clark

Principal-Agent Problems and Experimentation

Abstract

Chapter 1 studies contracts which provide incentives for information acquisition that is costly for an agent and unobservable to a principal. We show that each Pareto optimal contract pays a fraction of output, a state-dependent transfer, and a decision-dependent transfer. The fraction of output is increasing in the set of experiments available to the agent, the state-dependent transfer indexes contract payments to differences in output between states, and the decision-dependent transfer punishes decisions that are likely to make liability limits bind.

Chapter 2 studies the effect of capacity constraints capping maximum agent effort in principal-agent problems. We show that the form of optimal contracts after introducing a capacity constraint can be obtained from the form of the original optimal contracts by replacing output with a fraction of output and we show that this fraction is increasing in the agent's capacity. One implication of this result is that if before introducing a capacity constraint all optimal contracts result in the principal receiving a debt, then optimal contracts after introducing a capacity constraint result in the principal receiving a debt and a fraction of equity.

Chapter 3 studies the relationship between core and competitive equilibria in economies that consist of a continuum of agents and some large agents. We construct a class of these economies in which the core and competitive allocations do not coincide.

Contents

Abstract	iii
Acknowledgments	ix
1 Contracts for Experimentation	1
1.1 Introduction	1
1.1.1 Background	4
1.2 Framework	6
1.3 Analysis and results	9
1.3.1 First-best contracts	10
1.3.2 Dealing with the capacity constraint	11
1.3.3 Second-best contracts	12
1.3.4 Piece rates, indexation, and the optimal distortion	14
1.3.5 Posterior-separable costs	15
1.3.6 Multiple periods	18
1.4 Discussion	20
1.4.1 Debt and equity	20
1.4.2 Restricted investment contracts	21
1.4.3 Parallel literatures	23
1.5 Summary	24
2 Capacity Constraints in Principal-Agent Problems	25
2.1 Introduction	25
2.2 Framework	27
2.3 Examples	28
2.4 Result	29
2.5 Proof	31
2.6 Application to a theory of capital structure	32
2.6.1 Live-or-Die Contracts	33

2.7	Discussion	33
2.8	Conclusion	35
3	Core Equivalence with Large Agents	36
3.1	Framework	37
3.2	Analysis and results	41
3.2.1	A rescaling lemma	43
3.2.2	A two-atom economy	45
3.2.3	Constructing the core allocation	47
3.2.4	Proof of the main theorem	49
3.3	Graphical proof for the two commodity case	50
	References	53
	Appendix A Appendix to Chapter 1	58
A.1	Proofs	58
A.1.1	Proof of Proposition 1	58
A.1.2	Proof of Proposition 2	59
A.1.3	Proof of Proposition 3	60
A.1.4	Results used in the proofs of Propositions 1 and 3	64
A.1.5	Derivation of the optimal contract in the multi-period model	66
A.1.6	Posterior-separable costs and complementarities	67
A.1.7	Risk-aversion	71
A.2	Graphical method and an example	72
A.2.1	Graphical method	72
A.2.2	A two-state two-decision example	77
	Appendix B Appendix to Chapter 2	87
B.1	Risk-aversion	87
B.2	Discounting	89
	Appendix C Appendix to Chapter 3	91
C.1	Perfect competition	91
C.1.1	The Arrow-Debreu economy	91
C.1.2	Debreu and Scarf: A limit theorem on the core of an economy	95
C.1.3	A model of perfect competition	102
C.1.4	Aumann: Markets with a continuum of traders	106

C.2	Imperfect Competition	112
C.2.1	Imperfect competition	113
C.2.2	A model of imperfect competition	113
C.2.3	Shitovitz: Oligopoly in markets with a continuum of traders	116
C.2.4	Gabszewicz and Mertens: Atoms that aren't "too" big	122

List of Tables

A.1	(a) First best experiment when the agent's capacity constraint does not bind. (b) First best experiment under capacity constraint $c(p) \leq 1/2$	78
A.2	(a) Optimal contract (b) Optimal experiment (c) State-dependent transfer (d) Optimal distortion	81
A.3	(a) Contract $y - \beta$. (b) First best experiment under contract $y - \beta$	81
A.4	Optimal experiment for contract $\max\{0, y - \beta\}$	82
A.5	Comparison of expected payoffs and costs	85

List of Figures

3.1	Proof of the main theorem: configuration 1.	51
3.2	Proof of the main theorem: configuration 2.	51
3.3	Proof of the main theorem: configuration 3. Interchange a_1 and a_2	52
A.1	The agent's problem for $\pi(\theta_2) = 0.45$ and contract $b(d_1, \theta_1) = 0, b(d_2, \theta_1) = b(d_2, \theta_2) = 1$, and $b(d_1, \theta_2) = 2$. Reduced form $B(p)$ and upper envelope of net utilities $B(p) + Y(p)$ as a function of the posterior probability assigned to state θ_2 . Optimal posteriors are given by maximizing expected net utility subject to the average posterior equaling the prior.	73
A.2	Altering the contract b by the state contingent payment $\beta(\theta_1) = 0, \beta(\theta_2) = 1$ does not alter the experiment the agent chooses.	75
A.3	(a) A binding capacity constraint (thick dashed lines) brings the posterior probabilities closer to the prior. (b) With a binding capacity constraint (thick dashed lines), scaling incentives in half (thick solid lines) does not alter the agent's chosen experiment.	76
A.4	Contract $y - \beta$ induces first best experimentation.	82
A.5	Optimal experiments for $\max\{0, y - \beta\}$ (thick lines) and $y - \beta$ (thin lines).	83
A.6	Point A gives expected output minus $E_\pi[\beta(\theta)]$. Point B gives the expected contract payment.	84
A.7	The optimal contract. The lines defining A and B coincide so the optimal contract $b = y - \beta - \gamma$ has the same expected value as the contract $y - \beta$ at the optimal posteriors.	85
A.8	A perturbation (thick line) away from the optimal contract. The expected contract payment (middle dot) falls less than expected output (lower dot).	86
C.1	Forming the set Γ	99
C.2	Forming a blocking coalition with Lyapunov's convexity theorem	109
C.3	From imperfect to perfect competition	118
C.4	Forming a blocking coalition	126

Acknowledgments

I am extremely grateful to my advisors Oliver Hart, Mihai Manea, and Eric Maskin. Each has been immensely generous with their time and advice. I also wish to thank Adrien Auclert, Kirill Borusyak, Gabriel Carroll, Ben Golub, Tommaso Denti, Richard Holden, Bengt Holmström, Ryota Iijima, Michihiro Kandori, Jonathan Libgober, Xiaosheng Mu, Jeff Picel, Harry Di Pei, Matt Rognlie, Andrei Shleifer, Tomasz Strzalecki, Yao Zeng, and audiences at Harvard and MIT for helpful comments on Chapters 1 and 2. I thank Harvard University for generous funding and the Colditz family for their hospitality during the summer of 2016. I cannot fully acknowledge here the great debt I owe to those I have worked alongside during the PhD, but I cannot refrain from singling out several key contributors. I am extremely grateful to Brenda Piquet, the administrator of the graduate program, for her sound advice and support. I thank Drew Fudenberg for instruction and guidance during my early years in graduate school. Finally, I thank my classmates: David Baquee, Tilman Dette, Mike Egedal, Peter Ganong, Duncan Gilchrist, Michael Gill, Tom Gole, Ben Hébert, Simon Jäeger, Rohan Kekre, Divya Kirti, Akos Lada, Annie Liang, Morgan McLellon, Vladimir Mukharlyamov, Joey Neggens, Alex Peysakovich, Mikkel Plagborg-Møller, Pepe Montiel Olea, Danial Pollman, Guillaume Pouliot, Tom Powers, Giovanni Reggiani, Gleb Romanyuk, Martin Rotemberg, Jann Spiess, Eduard Talamàs, Jing Xia, Danny Yagan, Fernando Yu, and John Zhou.

Chapter 1

Contracts for Experimentation

1.1 Introduction

Consider a principal that wishes to hire an agent to acquire information and take a decision on his behalf. The nature of the information precludes verifying how carefully the agent carries out this task, but the pair can enforce a contract stipulating transfers as a function of the decision the agent takes and the realized state of the world. Information is costly for the agent and he is capacity constrained in the amount he can acquire. Both parties are risk-neutral and transfers between them are constrained by liability limits.

We show that each Pareto optimal contract pays a fraction of output, a state-dependent transfer, and—with information costs proportional to expected Shannon entropy reduction—a decision dependent transfer. The fraction of output is increasing in the set of experiments available to the agent, the state-dependent transfer indexes contract payments to differences in output between states, and the decision-dependent transfer punishes decisions likely to make liability limits bind.

The resulting contract can be described simply: stipulation of the liability limits,

a punishment or reward for each decision based on its ex-ante risk of causing bankruptcy, an indexing payment so the agent is not rewarded purely for finding himself in a lucky state, and a piece rate based on the agent's capacity for acquiring information.

For a general cost function the expression that becomes decision-dependent under expected Shannon entropy reduction sheds light on how optimal incentives are shaped by complementarities in the cost of acquiring information. If a contract payment is stuck against a liability limit, then other contract payments substitute for its role of influencing the signal probability; this occurs by exploiting complementarities in the cost of acquiring different signals in different states.

We do not impose restrictions on contract form, such as monotonicity or linearity, or on the way the agent can influence the distribution over output, such as monotone likelihood ratio or distribution function convexity properties. Our model is at its core an information acquisition version of the static model in Holmström and Milgrom's classic paper [33] on linear contracts, in which the domain of the agent's cost function is a feasible set of distributions, rather than a set of effort parameters. This side-steps the need to specify a mapping between effort parameters and the distribution over output.

This approach also starts the analysis in the right place for thinking about linear contracts. A theme in that literature is that linear contracts arise because they are robust to uncertainty about the contracting environment and to freedom in the way the agent can influence the distribution over output (e.g. [12, 22]). To see the second of these points, consider the following argument for why linear contracts align the incentives of a principal and agent: where it can be avoided non-linearity is wasteful because it drives apart the agent's and principal's preferred distributions over output; a convex payoff for the agent nearby some output level means a concave

payoff for the principal, and thus an agent who wishes for higher variation in output than the principal nearby this level of output. This argument is strengthened by the agent having freedom in how he can influence the distribution over output. We take advantage of this by starting out with a model that allows the agent to choose any distribution over output. This starting point also forces us to think carefully about the restrictions we do place on the agent's feasible set, such as the capacity constraint.

Our results demonstrate the importance of capacity constraints in principal-agent models by showing that they pin down the fraction of output paid to the agent. In the context of modeling information acquisition, usual practice dictates that a decision maker either pays a cost for information in terms of utility, or faces a budget set of feasible experiments given by a capacity constraint. These approaches are interchangeable in decision problems—one is dual to the other. However, in principal-agent models the optimal contract varies significantly depending on whether or not a capacity constraint is imposed. The results we obtain regarding capacity constraints do not depend on the agent choosing an experiment. Rather, they apply to any standard principal-agent setup with a risk-neutral principal and agent. For example, one could augment the security design model of Innes [35] with a capacity constraint and apply our results. The conclusion of that paper in which debt is optimal then changes, the optimal contract now being debt and a fraction of equity. This connects our work to the literature that seeks to derive optimal capital structure from agency considerations (e.g. [36]).

The rest of the paper is organized as follows. We lay out the model in section 1.2. Section 1.3 contains the analysis and results, and is divided into six parts: in section 1.3.1 we characterize the Pareto optimal contracts that are first-best; in section 1.3.2 we present a method for dealing with the capacity constraint, allowing

subsequent analysis to do without it; section 1.3.3 studies the form of Pareto optimal contracts that are second-best; section 1.3.4 discusses how the previous three sections work together to characterize the set of Pareto optimal contracts; in section 1.3.5 we consider the contract forms arising from different cost functions and discuss the simplest way of describing the optimal contract; and in section 1.3.6 we present a multi-period model. We discuss applications and relations to other contract forms in the literature in section 1.4; in section 1.5 we conclude. All proofs and derivations are collected in appendix A.1, and appendix A.2 solves a simple example and presents a graphical method for analyzing the optimal contract.

1.1.1 Background

The first study of optimal incentives for information acquisition was completed by Demski and Sappington [21]. They focused on optimal risk-sharing arrangements and on contracts that depend solely on output. Their main concern was the interaction between the “planning phase” (the acquiring of information) and the “implementation phase” (the choice of a decision). Their basic point was that, even if implementation is costless (as it is in our model), the unverifiable nature of planning induces moral hazard in implementation since the outcome of the agent’s decision provides information about his planning.

Osband [46] was led to the problem from his study of proper scoring rules—incentive schemes that elicit true beliefs (see [49]). His goal was to characterize the best proper scoring rule eliciting a forecaster’s estimate of the mean of a random variable. The forecaster draws observations of the random variable at a cost, and this cost as well as the number of observations the forecaster makes are unknown to the principal. His focus was on the merits of (a) screening by offering a menu of contracts versus (b) pitting forecasters in competition against one another.

More recently, Zermeño [59, 60] has developed a very general model of incentivizing information acquisition under liability limits. His analysis focuses on the interaction between the verifiable data on which a contract can depend (for instance, the decision may be verifiable but the state may not be) and: (1) the usefulness of menus of contracts, and (2) whether the principal or agent is tasked to take the decision. Menus are useful because they allow contract payments to partially depend on the state: following receipt of a signal the agent gets to choose a decision and a contract from the menu; since the signal is correlated with the state the menu functions as a tool for correlating contract payments with the state—a role that is otherwise forced upon the agent’s choice of decision—with a more precise signal leading to better correlation. In general, menus do not allow perfect correlation between contract payments and the state, and for this reason decision making is usually ex-post inefficient—it is distorted so that the agent is more likely to take a decision that reveals the state. Regarding (2), the allocation of decision making authority is irrelevant when the principal and agent can infer from the verifiable data which decision was taken; when this is unclear, whether the principal or agent is tasked to take the decision affects the set of implementable outcomes.

Carroll [13] builds on and refines this model to study the form of “robust” incentives for information acquisition: a principal does not know all experiments or experiment costs available to an agent and ranks contracts according to their minimum expected payoff among all experiments and experiment costs including a known set. He shows that the optimal contract is a *restricted investment contract*: the set of decisions available to the agent is restricted, and for unrestricted decisions the optimal contract pays a fraction of output and a state-dependent transfer.

Our paper falls closest to Carroll’s. In the terminology of our model, one may view Carroll’s restriction on decisions as resulting from a decision-dependent trans-

fer that severely punishes restricted decisions. From this perspective, the contracts we find generalize restricted investment contracts. This provides a link between Bayesian and non-Bayesian approaches to designing incentives for experimentation.

1.2 Framework

There are two dates. At date 0, a principal faces a choice from a finite set of *decisions* D . The date 1 payoff to the principal from a decision depends on the realization of one among a finite set of *states* Θ , as described by an *output function* $y : D \times \Theta \rightarrow \mathbb{R}$. The principal hires an agent to acquire information about the state and take a decision on his behalf but cannot either monitor how much care the agent takes acquiring information or observe the information the agent does acquire. The principal motivates the agent with a *contract* $b : D \times \Theta \rightarrow \mathbb{R}$ that depends on the decision the agent takes and the realized state of the world—both are verifiable.

At date 0, the principal and agent agree to a contract b . At date 1, before choosing a decision, the agent performs an *experiment* $(X, \{p(\cdot|\theta)\}_{\theta \in \Theta})$, which consists of a finite set of *signals* X and a collection of probability distributions $\{p(\cdot|\theta)\}_{\theta \in \Theta} \subseteq \Delta(X)$ on the set of signals. A *decision rule* $f : X \rightarrow \Delta(D)$ describes the agent's randomization over decisions upon observing each signal.

An *action profile* $(b, X, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$ consists of a contract, experiment, and decision rule. The agent has preferences over action profiles $(b, X, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$ given by the expected value of contract payments less an experiment cost:

$$\sum_{\theta \in \Theta} \pi(\theta) \sum_{x \in X} p(x|\theta) \sum_{d \in D} f(d|x) b(d, \theta) - c(X, \{p(\cdot|\theta)\}_{\theta \in \Theta}).$$

The distribution $\pi \in \Delta(\Theta)$ is the agent's *prior belief* over states and the function c is the agent's *cost function*. For now, the cost function is only assumed to respect the

ordering of experiments given by informativeness (see Blackwell's theorem¹).

The principal has preferences over action profiles $(b, X, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$ given by the expected value of output less contract payments:

$$\sum_{\theta \in \Theta} \pi(\theta) \sum_{x \in X} p(x|\theta) \sum_{d \in D} f(d|x) (y(d, \theta) - b(d, \theta)).$$

Contracts are restricted to belong to a set \mathcal{B} of *feasible contracts* and experiments are restricted to belong to a set \mathcal{E} of *feasible experiments*.

A contract b satisfies the *limited liability* constraints if $0 \leq b(d, \theta) \leq y(d, \theta)$ for all $d \in D, \theta \in \Theta$. The set of feasible contracts are those satisfying the limited liability constraints. This may mean that the principal and agent are wealth constrained (with wealth normalized to zero or included in output y) or that the contract is subject to statutory liability limits (as is the case for a limited liability corporation).

In addition to the incremental cost of acquiring information represented by the cost function c , we assume that the agent faces a *capacity constraint* $c(X, \{p(\cdot|\theta)\}_{\theta \in \Theta}) \leq k$ that limits the experiments he can perform. The set of feasible experiments are those satisfying this constraint.

An action profile $(b, X, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$ is *feasible* if the contract b and experiment

¹Experiment $(X, \{p(\cdot|\theta)\}_{\theta \in \Theta})$ is *more informative* than experiment $(X', \{p'(\cdot|\theta)\}_{\theta \in \Theta})$ if for any prior π and contract b , whenever some expected payoff can be achieved under the latter experiment, then it can be achieved under the former. That is, for each decision rule f' on X' there exists a decision rule f on X such that

$$\sum_{\theta \in \Theta} \pi(\theta) \sum_{x' \in X'} p'(x'|\theta) \sum_{d \in D} f'(d|x') b(d, \theta) = \sum_{\theta \in \Theta} \pi(\theta) \sum_{x \in X} p(x|\theta) \sum_{d \in D} f(d|x) b(d, \theta).$$

Experiment $(X, \{p(\cdot|\theta)\}_{\theta \in \Theta})$ is *sufficient* for experiment $(X', \{p'(\cdot|\theta)\}_{\theta \in \Theta})$ if there is a *garbling function* $g : X \rightarrow \Delta(X')$ such that

$$p'(x'|\theta) = \sum_{x \in X} p(x|\theta) g(x'|x).$$

That is, experiment $(X', \{p'(\cdot|\theta)\}_{\theta \in \Theta})$ can be performed by randomizing over its outcomes according to g where the randomization depends on the outcome of the experiment $(X, \{p(\cdot|\theta)\}_{\theta \in \Theta})$.

Blackwell's theorem [8] states that an experiment is more informative than another if and only if it is sufficient.

$(X, \{p(\cdot|\theta)\}_{\theta \in \Theta})$ are feasible, and there is no action profile $(b, X', \{p'(\cdot|\theta)\}_{\theta \in \Theta}, f')$, consisting of the same contract and some feasible experiment, that the agent prefers.

An action profile can be *improved upon* if there is a feasible profile that makes either the agent or principal better off without making the other worse off. An action profile is *Pareto optimal* if it is feasible and cannot be improved upon.

Normalizing the experiment and decision rule. A decision rule is a garbling function (see footnote 1). Thus, given action profile $(b, X, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$, we may define a garbled experiment $(D, \{p'(\cdot|\theta)\}_{\theta \in \Theta})$ by $p'(d|\theta) = \sum_{x \in X} p(x|\theta)f(d|x)$ and choose the decision rule $f' : D \rightarrow \Delta(D)$ that maps each decision to the degenerate distribution on that decision. The contract has the same expected payoff under this experiment and decision rule and since the former experiment is sufficient for the latter it is weakly more costly. Thus, it is without loss to set the agent's set of signals to be D and the agent's decision rule $f : D \rightarrow \Delta(D)$ to map each decision to the degenerate distribution on that decision.

Given this normalization, we will write p for the experiment $(D, \{p(\cdot|\theta)\}_{\theta \in \Theta})$, and write (b, p) for the action profile $(b, D, \{p(\cdot|\theta)\}_{\theta \in \Theta}, f)$ in which f maps each decision to the degenerate distribution on that decision. We will write $E_p[b] - c(p)$ and $E_p[y - b]$ for the agent's and principal's utility from action profile (b, p) . For experiment p , we will write $p(d, \theta)$ for the joint probability $p(d|\theta)\pi(\theta)$, $p(d)$ for the total probability $\sum_{\theta} \pi(\theta)p(d|\theta)$, and $p(\theta|d)$ for the posterior probability $p(d|\theta)\pi(\theta) / \sum_{\theta'} \pi(\theta')p(d|\theta')$.

1.3 Analysis and results

Since both principal and agent are risk-neutral, a Pareto optimal contract is one that maximizes welfare—expected output less the cost of the agent’s chosen experiment—subject to the agent obtaining a given level of utility. A way to characterize all Pareto optimal contracts is to start with the contract equal to output and to then consider feasible alterations that increase the principal’s payoff at minimal loss to welfare. Our setup provides two ways to do this without altering the agent’s behavior, and thus without any loss in welfare.

First, since the agent is risk-neutral and does not control the state, altering contract payments by a state-dependent transfer does not alter his behavior. Second, if the agent’s capacity constraint binds then a slight scaling down of the contract—multiplying it by a number slightly less than one—will not alter the agent’s behavior provided the capacity constraint continues to bind.

The interval of utilities achievable for the agent by applying these maneuvers to output, correspond to the points on the Pareto frontier in which welfare is maximized. Proposition 1 characterizes these contracts. For lower agent utilities there is no way to increase the principal’s payoff without reducing welfare, and thus for such contracts the agent’s liability limit binds for at least one decision in each state and the agent’s capacity constraint does not bind.

We characterize these second-best Pareto optimal contracts in Proposition 3. These contracts share the features of first-best contracts in that they pay a fraction of output less a state-dependent transfer. But they also involve an optimal distortion, which we give a formula for in this proposition.

In section 1.3.5, we provide a simplified version of this formula specialized to the class of posterior separable cost functions studied in [10] and [31]. Expected

reduction in Shannon entropy belongs to this class. Under this cost function the optimal distortion reduces to a decision-dependent transfer in which the agent is punished for taking decisions that are likely to make his liability limits bind and rewarded for decisions that are likely to make the principal's liability limits bind.

Finally, in section 1.3.6 we study, with a multi-period version of our model, how the optimal contract changes when the agent's decision and the state control future, as well as current, output. The optimal contract has the same form as in the static model except now the agent is led to internalize the effect of his decision on the principal's future payoffs: the results of the static model apply with current output replaced by output together with the principal's continuation payoff.

1.3.1 First-best contracts

In the standard principal-agent setup with risk-neutral parties the *first-best* contracts (those that maximize welfare $E_p[y] - c(p)$) are given by output less a constant, $b = y - t$. If the principal "sells the firm" to the agent for a fee t , then the agent necessarily maximizes welfare. In problems of information acquisition there are additional first-best contracts. Given that the agent is risk-neutral and does not control the state, modifying a contract by a state-dependent transfer does not alter his incentives. Thus, any $y - \beta$ with $\beta : \Theta \rightarrow \mathbb{R}$ is first-best.

Introducing a capacity constraint into the problem introduces more first-best contracts. There is $\alpha' \in [0, 1]$ such that for any first-best contract $y - \beta$ all contracts $\alpha y - \beta$ with $\alpha \in [\alpha', 1]$ are also first-best. When the agent's capacity constraint does not bind under contract $y - \beta$, then $\alpha' = 1$; if it does bind, α' is less than 1.

Conversely, if the cost function, viewed as a function on $\mathbb{R}^{|D \times \Theta|}$, is differentiable, and each first-best experiment assigns positive probability to each decision in each state, then all first-best contracts take this form.

Proposition 1. Define α' by

$$\alpha' = \sup \{ \alpha \in [0, 1] : c(p) < k \text{ for all feasible } p \text{ optimal for agent against contract } \alpha y \}.$$

If a contract has the form $\alpha y - \beta$, with $\alpha \geq \alpha'$ and $\beta : \Theta \rightarrow \mathbb{R}$, then it is first-best.

Conversely, suppose the cost function is differentiable and all first-best experiments assign positive probability to each decision in each state. Then, if b is a first best contract, $b = \alpha y - \beta$ for some $\alpha \geq \alpha'$ and $\beta : \Theta \rightarrow \mathbb{R}$.

Note that if a first-best contract is feasible, it is also Pareto optimal. Notice too that if there is a first-best Pareto optimal contract yielding the agent utility r , then all Pareto optimal contracts yielding the agent utility at least r are first-best since the agent's utility can be increased without affecting his behavior.

From here, it is easy to see that the minimum agent utility achievable with a first-best Pareto optimal contract is obtained by the contract $\alpha' y - \beta$, where $\beta : \Theta \rightarrow \mathbb{R}$ brings the minimum contract payment in each state to zero: $\beta(\theta) = \min\{\alpha' y(d, \theta) : d \in D\}$.

Also note that α' is increasing in the agent's capacity k .

1.3.2 Dealing with the capacity constraint

In this section we show that imposing a capacity constraint is equivalent to scaling output. This will allow us to proceed in solving for the second-best Pareto optimal contracts with a scaled level of output in place of the capacity constraint.

The key idea is to consider a perturbation of our model, parametrized by $\alpha \in [0, 1]$, in which output y in the principal's objective function is replaced by αy (the feasible contracts, distributions, and profiles remain unchanged). Denote the set of Pareto optimal profiles to the perturbed problem by $\mathcal{P}(\alpha)$, and denote by

$\mathcal{P}(\alpha, r)$ the subset of $\mathcal{P}(\alpha)$ in which the agent's utility is minimized subject to it being at least equal to r .

Consider a Pareto optimal profile $(b, p) \in \mathcal{P}(1, r)$ to the unperturbed problem. The following Proposition shows that all solutions $\mathcal{P}(1, r)$ of this unperturbed problem (a problem in which the capacity constraint might bind) may be obtained as solutions $\mathcal{P}(\alpha^*, E_p[b] - c(p))$ to the perturbed problem for the $\alpha^* \in [0, 1]$ defined in the proposition and that, if the agent's cost function is continuous, then for at least one solution of this problem the capacity constraint does not bind.

Proposition 2. *Let $(b, p) \in \mathcal{P}(1, r)$. Then there exists $\alpha^* \in [0, 1]$, defined as*

$$\alpha^* = \sup \{ \alpha \in [0, 1] : c(p') < k \text{ for all } (b', p') \in \mathcal{P}(\alpha, E_p[b] - c(p)) \},$$

so that for each $\alpha \in [\alpha^*, 1]$,

$$\mathcal{P}(\alpha, E_p[b] - c(p)) \supseteq \mathcal{P}(1, r),$$

and, conversely, if $(b_\alpha, p_\alpha) \in \mathcal{P}(\alpha, E_p[b] - c(p))$ is such that $c(p_\alpha) = k$, then

$$(b_\alpha, p_\alpha) \in \mathcal{P}(1, r).$$

If the agent's cost function is continuous, then there is a Pareto optimal profile in $\mathcal{P}(\alpha^, E_p[b] - c(p))$ that solves the problem obtained by removing the capacity constraint from this perturbed problem.*

Note that α^* , by its definition, is increasing in the agent's capacity.

1.3.3 Second-best contracts

Having dealt with first-best Pareto optimal contracts (and developed a tool to jettison the capacity constraint) we can now turn to their second-best counterparts.

Proposition 2 tells us that we can do away with the capacity constraint and instead replace output by α^*y . In the next proposition we solve for the form of the optimal contract in the resulting problem.

Proposition 3. *Let (b, p) belong to $\mathcal{P}(\alpha, r)$ and assume that*

- *the agent's capacity constraint does not bind: p is an optimal experiment for the agent without a capacity constraint when the contract is b ;*
- *the agent's cost function c is strictly convex and its second derivative is continuous;*
- *p assigns positive probability to each decision in each state.*

Then

$$b(d, \theta) = \alpha y(d, \theta) - \beta(\theta) - \gamma(d, \theta),$$

$$\gamma(d, \theta) = \frac{1}{\pi(\theta)} \sum_{d', \theta'} \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')} \left(p(d'|\theta') (1 - \xi) - \frac{\lambda[d', \theta']}{\pi(\theta')} \right). \quad (1.1)$$

The term $\lambda[d, \theta]$ is a Lagrange multiplier for the constraint $0 \leq b(d, \theta) \leq y(d, \theta)$: it is non-negative when only the agent's liability limit binds, non-positive when only the principal's liability limit binds, zero when neither liability limit binds, and unrestricted in its value if both liability limits bind; $\xi \in [0, 1]$ and is a Lagrange multiplier on the agent's participation constraint.

By Proposition 2, when the agent is capacity constrained there is a Pareto optimal contract of the form

$$b(d, \theta) = \alpha^* y(d, \theta) - \beta(\theta) - \gamma(d, \theta),$$

where $\alpha^* \in [0, 1]$ and is increasing in the agent's capacity. For second best contracts we have that

$$\beta(\theta) = \min\{\alpha^* y(d, \theta) - \gamma(d, \theta) : d \in D\}.$$

Note that if the mapping from $\alpha \in [0, 1]$ to the set of Pareto optimal profiles $\mathcal{P}(\alpha, r)$ is lower hemi-continuous at α^* , then all second-best Pareto optimal contracts take this form (cf. Proposition 2).

1.3.4 Piece rates, indexation, and the optimal distortion

Together, propositions 1, 2, and 3 provide a general characterization of Pareto optimal contracts in terms of a piece rate $\alpha \in [0, 1]$, a state-dependent transfer $\beta : \Theta \rightarrow \mathbb{R}$, and an optimal distortion $\gamma : D \times \Theta \rightarrow \mathbb{R}$.

For a first-best contract, $\gamma = 0$ and there is no distortion to incentives. The piece rate α may be any value in the interval $[\alpha', 1]$, α' being the point below which the agent's capacity is no longer exhausted. The state-dependent transfer β is any state dependent transfer such that the resulting contract $\alpha y - \beta$ satisfies the liability limits and so is feasible.

For a second-best contract the piece rate is equal to α^* and the state-dependent transfer β is given by $\beta(\theta) = \min\{\alpha^* y(d, \theta) - \gamma(d, \theta) : d \in D\}$ so that liability limits bind in each state. The non-distortionary ways of transferring utility from the agent to the principal are fully exploited.

The term γ describes the optimal distortion to incentives—that is, it is the way of transferring utility from the agent to the principal at minimum loss to welfare (expected output less experiment cost).

Expression 1.1 sheds light on how this optimal distortion is constructed showing that it largely depends on the complementarities in the cost of acquiring different signals in different states. To see this, suppose that the agent's liability limit binds in state θ' following decision d' so that the Lagrange multiplier $\lambda[d', \theta'] > 0$ (i.e. the optimal contract would be different without this liability limit). To lever down the probability of receiving the signal to take decision d' when the state is

θ' the principal can no longer reduce the contract payment $b(d', \theta')$. As a result, other contract payments—those not immobilized against their liability limits—are recruited for this task. Expression 1.1 shows the way this happens. If a contract payment $b(d, \theta)$ is not stuck against the agent's liability limit then complementarity in the cost of signal probabilities $p(d'|\theta')$ and $p(d|\theta)$ together with $\lambda[d', \theta'] > 0$ adds to $\gamma(d, \theta)$. This reduces the contract payment $b(d, \theta)$ which in turn reduces the signal probability $p(d|\theta)$ and hence signal probability $p(d'|\theta')$.

The term $p(d'|\theta')(1 - \xi)$ in Expression 1.1 accounts for adjustments in the contract payments $b(d, \theta)$ necessary to keep the incentive and participation constraints holding following a reduction in $p(d'|\theta')$.

1.3.5 Posterior-separable costs

Complementarities in information acquisition

Here we consider a class of cost functions defined in [31]. These cost functions emerge from a dynamic information acquisition problem in which an agent at each point in time decides whether to take a decision or continue collecting information. Information collection is modeled as the agent choosing the covariances of a multi-dimensional Brownian motion that moves about the set of probability distributions over states, subject to a constraint on these covariances in terms of the complementarities and substitutabilities of acquiring information about different states. They show that this dynamic problem has a static representation in which the agent's cost function is defined as

$$c(p) = \sum_{d \in D} p(d) D(p(\cdot|d) || \pi)$$

where

$$D(p(\cdot|d)||\pi) = H(p(\cdot|d)) - H(\pi) - (p(\cdot|d) - \pi)^T \nabla H(\pi)$$

is the Bregman divergence associated with the convex function H and where

$$\frac{\partial}{\partial p(\theta|d)} \frac{\partial}{\partial p(\theta'|d)} H(p(\cdot|d)) = \frac{k(\theta, \theta', p(\cdot|d))}{p(\theta|d)p(\theta'|d)}. \quad (1.2)$$

The matrix k is called the information cost matrix and characterizes the complementarities and substitutabilities in acquiring information about different states. It is positive semi-definite, symmetric, and its rows sum to zero. Given experiment p , there is complementarity in learning about distinct states θ and θ' when $k(\theta, \theta', p(\cdot|d))$ is negative and substitutabilities when it is positive; when $\theta = \theta'$ it is nonnegative and measures the difficulty of learning about state θ .

For any cost function in this class, its Hessian is given by (see section A.1.6 for the derivation)

$$\frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d|\theta')} = p(d) \frac{k(\theta, \theta', p(\cdot|d))}{p(d|\theta)p(d|\theta')},$$

and 0 otherwise. Substitution into (1.1), the expression for γ , gives

$$\gamma(d, \theta) = -\frac{1}{\pi(\theta)} \sum_{\theta'} \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d|\theta')} \frac{\lambda[d, \theta']}{\pi(\theta')}.$$

A simple description of optimal contracts

A simple and intuitive form for γ arises when the agent's information cost matrix k is given by the inverse Fisher information matrix

$$k(\theta, \theta', p(\cdot|d)) = \begin{cases} p(\theta|d)(1 - p(\theta|d)) & \text{if } \theta = \theta' \\ -p(\theta|d)p(\theta'|d) & \text{if } \theta \neq \theta' \end{cases}.$$

Then the agent's cost function is expected reduction in Shannon entropy

$$c(p) = H_S(\pi) - \sum_{d \in D} p(d) H_S(p(\cdot|d)),$$

where H_S denotes Shannon entropy, defined for each $q \in \Delta(\Theta)$ by

$$H_S(q) = - \sum_{\theta \in \Theta} q(\theta) \log q(\theta).$$

Then

$$\gamma(d, \theta) = \sum_{\bar{\theta}} \frac{\lambda[d, \bar{\theta}]}{p(d)} - \frac{\lambda[d, \theta]}{p(d|\theta)\pi(\theta)}$$

and so the optimal contract has the following simple description:

- **Decision d penalty/reward:** $\hat{\gamma}(d) = \sum_{\bar{\theta}} \frac{\lambda[d, \bar{\theta}]}{p(d)}$;
- **Indexed output:** $y_I = y - \beta/\alpha^*$ where $\beta(\theta) = \min\{\alpha^* y(d, \theta) - \hat{\gamma}(d) : d \in D\}$;
- **Piece rate:** $\alpha^* \in [0, 1]$.

Here we assume the liability limits are implicitly understood: if the sum of promised payments falls outside the allowed range, the actual payment is the closest liability limit.

Notice that the decision d penalty/reward depends on the number of states in which liability limits bind following decision d . When the agent's liability limit binds $\lambda[d, \theta] > 0$ which pushes the decision d transfer in the direction of a punishment; when the principal's liability limit binds $\lambda[d, \theta] < 0$ which pushes the decision d transfer in the direction of a reward. The magnitude of the multiplier indicates the cost to the principal of having the liability limit bind.

One can see from Expression (1.2) that we recover the same contract form as when the agent's cost function is expected reduction in Shannon entropy provided

the information cost matrix is of the form

$$k(\theta, \theta', p(\cdot|d)) = p(\theta|d)g(\theta', p(\cdot|d)) + \mathbf{1}_{\{\theta'=\theta\}}h(\theta, p(\cdot|d))$$

for some functions f and g . However, from this form it is straightforward to show (using symmetry of the information cost matrix and that its rows sum to zero) that the cost function is necessarily proportional to expected reduction in Shannon entropy.

1.3.6 Multiple periods

In this section we study a multi-period version of our model in which the static contracting problem arises each period. Output and contracts now depend on histories of decisions and states, and the prior is a distribution over sequences of states.

Formally: there are a finite number of periods $t = 1, 2, \dots, T$, a finite number of states Θ one realizing each period, and a finite number of decisions D one being chosen each period. Output $y : \bigcup_{t=1}^T D^t \times \Theta^t \rightarrow \mathbb{R}$ depends on histories of decisions and states $(d^t, \theta^t) \in D^t \times \Theta^t$; output $y(d^t, \theta^t)$ is delivered at time t immediately following history (d^t, θ^t) . A contract $b : \bigcup_{t=1}^T D^t \times \Theta^t \rightarrow \mathbb{R}$ follows the same conventions as output but represents payments to the agent. The set of feasible contracts \mathcal{B} are those satisfying the principal's and agent's liability limits $0 \leq b(d^t, \theta^t) \leq y(d^t, \theta^t)$ for all (d^t, θ^t) in $\bigcup_{t=1}^T D^t \times \Theta^t$. An experiment is now a set of signals X and for each history (d^{t-1}, θ^{t-1}) a collection of probability distributions $\{p(\cdot|d^{t-1}, \theta^{t-1}, \theta)\}_{\theta \in \Theta} \subseteq \Delta(X)$ over the signals. A *decision rule* $f : X \times \bigcup_{t=1}^{T-1} D^t \times \Theta^t \rightarrow \Delta(D)$ maps signals at each history (d^{t-1}, θ^{t-1}) to randomizations over decisions. It is clear that we can again normalize the experiment and decision rule by applying the arguments of section 1.2: p denotes the experiment and decision

rule $(D, \{p(\cdot|d^{t-1}, \theta^{t-1}, \theta)\}_{\theta \in \Theta}, f)$ in which f maps signals in D to the degenerate distribution on the corresponding decision.

The set of feasible experiments \mathcal{E} is the set of all experiments such that for all histories (d^{t-1}, θ^{t-1}) the cost of the period t experiment $p(\cdot|d^{t-1}, \theta^{t-1}, \cdot)$ is less than the agent's capacity at this history $k(d^{t-1}, \theta^{t-1}) \in \mathbb{R}$. We also allow the agent's cost function to depend on the history of decisions and states; thus the cost of the period t experiment following history (d^{t-1}, θ^{t-1}) will be denoted $c(p|d^{t-1}, \theta^{t-1})$.

The preferences of the principal and agent are given by the discounted sum of their per period utilities (as defined in the static model) discounted geometrically by discount factors $\delta_P, \delta_A \in [0, 1]$ and the prior is given by conditioning the belief $\pi \in \Delta(\Theta^T)$ on the past history of states (this belief could also be made to depend on the past history of decisions, in which case that would also be conditioned upon).

Let $v_P(d^t, \theta^t)$ denote the principal's continuation payoff at history (d^t, θ^t) and let $v_A(d^t, \theta^t)$ denote the agent's continuation payoff at history (d^t, θ^t) . The principal's and agent's utilities may be written recursively as $E_p[y + \delta_P v_P - b|d^{t-1}, \theta^{t-1}]$ and $E_p[b + \delta_A v_A|d^{t-1}, \theta^{t-1}] - c(p|d^{t-1}, \theta^{t-1})$.

We can now apply the reasoning of Proposition 1, 2, and 3 to the problem of designing the period t contract payoffs following history (d^{t-1}, θ^{t-1}) (see section A.1.5 for the derivation) to obtain:

$$b(d^t, \theta^t) = \alpha^*(d^{t-1}, \theta^{t-1})(y(d^t, \theta^t) + \delta_P v_P(d^t, \theta^t)) - \beta(d^{t-1}, \theta^{t-1}) - \gamma(d^t, \theta^t),$$

where $\alpha^*(d^{t-1}, \theta^{t-1})$ in $[0, 1]$ is defined as in Proposition 2 and

$$\gamma(d^t, \theta^t) = \frac{1}{\pi(\theta|\theta^{t-1})} \sum_{d', \theta'} \frac{\partial^2 c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta) \partial p(d'|d^{t-1}, \theta^{t-1}, \theta')} \phi[d^{t-1}, d', \theta^{t-1}, \theta'],$$

$$\phi[(d^{t-1}, d'), (\theta^{t-1}, \theta')] = p(d|d^{t-1}, \theta^{t-1}, \theta')(1 - \zeta[d^{t-1}, \theta^{t-1}]) - \frac{\lambda[d^t, \theta^{t-1}, \theta']}{\pi(\theta'|\theta^{t-1})}.$$

As in the static model, for a second-best contract the agent's liability limit binds in each state so that

$$\beta(d^{t-1}, \theta^t) = \min\{\alpha^*(d^{t-1}, \theta^{t-1})(y(d^t, \theta^t) + \delta_P v_P(d^t, \theta^t)) - \gamma(d^t, \theta^t) : d \in D\}.$$

The decision-dependent transfer punishes decisions that are likely to make the current period's liability limits bind; now, though, the liability limits might bind due to variations in the principal's continuation payoff. For second-best contracts, the state-dependent transfer again plays an indexing role so that the agent's liability limits bind in each state; the difference now is that it indexes output and the agent's continuation payoff. The fraction of output paid to the agent again reflects his capacity at each history.

1.4 Discussion

1.4.1 Debt and equity

Our model describes interaction between an entrepreneur and investor in which the investor's role is to provide funding and the entrepreneur's job is to acquire information and take a decision. In the simplest case, when the agent's information cost is proportional to expected Shannon entropy reduction, our model predicts that

$$\begin{aligned} & y(d, \theta) - b(d, \theta) \\ &= \min\{y(d, \theta), (\beta(\theta) + \hat{\gamma}(d))/\alpha^*\} + (1 - \alpha^*) \max\{0, y(d, \theta) - (\beta(\theta) + \hat{\gamma}(d))/\alpha^*\}. \end{aligned}$$

Thus, the payoffs of the project are split into three pieces. A debt of $(\beta(\theta) + \hat{\gamma}(d))/\alpha^*$ and a $1 - \alpha^* : \alpha^*$ split between investor and entrepreneur of equity $\max\{0, y(d, \theta) - (\beta(\theta) + \hat{\gamma}(d))/\alpha^*\}$.

Debt. The debt's face value is

$$\frac{\beta(\theta) + \hat{\gamma}(d)}{\alpha^*}.$$

The state-dependent transfer β corresponds to indexation of the debt (certain classes of bonds exhibit this feature, their coupon rate being a reference rate plus a quoted spread). The decision-dependent transfer $\hat{\gamma}$ corresponds to a debt covenant, automatically increasing debts following a course of action that risks bankruptcy.

Equity. Debt is repaid first. What remains is equity, which is split between the investor and entrepreneur. Fraction α^* remains inside the firm to motivate work; additional equity kept inside has no effect since capacity constraints bind, so it used outside the firm to raise capital.

Debt vs. equity. A capacity constrained entrepreneur first sells equity since it does not distort his incentives. After a certain amount of equity is issued his capacity constraint will stop binding. If, before this point, the investor receives his required return, entrepreneur effort is first-best. Debt is issued only if more capital is needed.

1.4.2 Restricted investment contracts

In [13], Carroll studies a principal-agent problem of information acquisition where the principal does not know all experiments or experiment costs available to the agent. His principal ranks contracts according to their minimum expected payoff among all experiments and experiment costs including a known set. He shows that the optimal contract is a restricted investment contract: the set of decisions available to the agent is restricted and for unrestricted decisions the agent is paid a fraction of output less a state-dependent transfer.

Carroll's result assumes that for each convex combination of decisions there is a decision that dominates the combination: for each $t \in [0, 1]$ and $d, d' \in D$ there is a decision d'' such that $y(d'', \theta) \geq ty(d, \theta) + (1 - t)y(d', \theta)$ for all $\theta \in \Theta$. With a finite number of decisions this assumption implies that there is a dominant decision: a decision d' such that $y(d', \theta) \geq y(d, \theta)$ for all $d \in D, \theta \in \Theta$. This eliminates the need for information acquisition. Carroll suggests that the set of decisions may be replaced by convex combinations of decisions. This would be problematic in our setting because a randomization over decisions is not verifiable (which for a restricted investment contract would be important to determine if the selected decision is restricted). Thus in Carroll's model the agent makes a recommendation and the principal is left to take the decision.

To compare our results to Carroll's, it is useful to think of his restriction on decisions as a payment $\gamma : D \times \Theta \rightarrow \mathbb{R}$ that sufficiently punishes restricted decisions that they are left unchosen and is zero for unrestricted decisions. Thus, Carroll obtains the contract $b = \alpha y - \beta - \gamma$ where $\alpha \in [0, 1]$, $\beta : \Theta \rightarrow \mathbb{R}$ is defined as $\beta(\theta) = \min\{\alpha y(d, \theta) - \gamma(d, \theta) : d \in D_{UR}\}$, and $\gamma : D \times \Theta \rightarrow \mathbb{R}$ is defined to be 0 on unrestricted decisions D_{UR} and such that $b(d, \theta) = 0$ for restricted decisions $D \setminus D_{UR}$. In contrast to this, we obtain the contract $b = \alpha y - \beta - \gamma$ where $\alpha = \alpha^* \in [0, 1]$ (as defined in Proposition 2), $\beta : \Theta \rightarrow \mathbb{R}$ is defined as $\beta(\theta) = \min\{\alpha y(d, \theta) - \gamma(d, \theta) : d \in D\}$, and $\gamma : D \times \Theta \rightarrow \mathbb{R}$ is given by

$$\gamma(d, \theta) = \frac{1}{\pi(\theta)} \sum_{d', \theta'} \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')} \left((1 - \xi)p(d'|\theta') - \frac{\lambda[d', \theta']}{\pi(\theta')} \right),$$

where $\lambda[d, \theta]$ and ξ are dual variables for the liability limits $0 \leq b(d, \theta) \leq y(d, \theta)$ and participation constraint. When our cost function is proportional to expected Shannon entropy reduction, $\gamma(d, \theta) = \sum_{\theta'} \lambda[d, \theta'] / p(d) - \lambda[d, \theta] / p(d|\theta) \pi(\theta)$ which conditional on liability limits not binding does not depend on the state.

In both sets of results γ punishes decisions that are likely to induce liability limits to bind. In Carroll's model, γ either rules out decisions or does nothing. In our model γ is less absolute: decisions are punished in proportion to how often and at what cost they induce binding liability limits. In both models the state-dependent transfers serve a normalizing role, setting the minimum contract payment in each state to zero.

Finally, the strength of the agent's incentives is measured by α in both Carroll's results and our own. Carroll does not discuss comparative statics for this parameter, but we are able to show in our model that it is increasing in the agent's capacity. It would be interesting to see whether a similar comparative static holds in Carroll's model, perhaps in terms of enlargements in the set of experiments and costs known to be available to the agent.

1.4.3 Parallel literatures

Acquiring information may be socially destructive. Effort spent may be wasteful, as in Hirshleifer [32] and Spence [53], or it may lead to trade-destroying information asymmetry, as in Akerlof [1] and Rothschild and Stiglitz [48]. These forces affect contract design. For example, in a model where an agent must report a project's outcome, Townsend [54] finds that debt contracts are optimal because they efficiently ration monitoring costs. Motivated by the second force, the same conclusion is drawn by Dang, Gorton, and Holmström [17] and Yang [57]: debt contracts are optimal to securitize state-dependent cash-flows because they avoid trade-destroying information asymmetry between the buyer and seller. Other examples include [15, 16, 58].

Recently, a theoretical literature has emerged that studies dynamic contracts that incentivize experimentation (e.g. [7, 26, 28, 29, 34, 40]). This literature often

uses results from the strategic experimentation literature (e.g. [9, 38]) but typically specializes to the case of a one or two armed bandit problem. An exception is Chassang [14] who considers a very general model of dynamic information acquisition. He shows that limited liability becomes approximately non-binding with a long enough time horizon and he uses this to characterize a class of contracts that work well in a variety of environments.

Sims [51, 52] popularized the use of expected reduction in Shannon entropy as a cost function to model bounded rationality; although its roots in economics were much earlier, for example in [2, 3], and before that in the statistical decision theory literature (e.g. [20, 39]). Recent efforts have aimed to derive behavioral foundations for different forms of costly information acquisition and to test popular cost functions using state-dependent stochastic choice data obtained from experiments (e.g. [10, 11, 31, 45, 56]).

1.5 Summary

We have studied contracting for the production of an experiment and provided necessary and sufficient conditions for optimal contracts to take an intuitive linear form. All Pareto optimal contracts pay a fraction of output, a state-dependent transfer, and—with information costs proportional to expected reduction in Shannon entropy—a decision-dependent transfer. For a general cost function, our characterization shows how complementarities in the cost of receiving different signals in different states affects the design of the optimal contract. In terms of these contracts, we have provided explanations for the strength of incentives, the role of indexing, and the role of punishments and rewards observed in actual contracts.

Chapter 2

Capacity Constraints in Principal-Agent Problems

2.1 Introduction

A widespread economic situation is the relation between a principal and an agent (see [4]). In the hidden action version of this problem an agent takes a costly action that affects the welfare of a principal. The relation between a physician and her patient, potential parties of a tort, investor and entrepreneur, and landlord and tenant are prominent examples. A feature of this relation is the set of actions available to the agent—the ways the agent can influence the distribution over outcomes. This set is determined by constraints on the environment (constraints on resources, technological know-how, etc). For example, a physician is constrained by the state of medical knowledge at the time she must make a diagnosis as well as by the amount of time and hospital resources she can devote to her patient. How should incentives change for the agent with only half the time to perform her task or after new technology expands the actions available to her?

This chapter addresses an aspect of this question. The agent's action is modeled as a distribution over an unknown state whose realization determines outcomes; attached to each distribution is the cost the agent incurs if he wishes to implement it; and the agent's choice is constrained to a feasible set of such distributions. We focus on modifying this feasible set via a capacity constraint capping the maximum cost of effort. Our main result applies when the principal and agent are risk-neutral. Here we obtain a relationship between optimal contracts with and without the capacity constraint.

To understand the result, imagine an optimal contract for which the agent's capacity constraint binds—that is, if the cap on effort cost were slightly relaxed the agent would change action. What relation does this contract bear to the optimal contract obtained after doing away with the capacity constraint? What we show is that imposing the capacity constraint is equivalent to scaling down output by some factor α^* in $[0, 1]$. That is, we find α^* such that given a mapping from output to optimal contracts for the problem without a capacity constraint, an optimal contract in the problem with a capacity constraint may be obtained by applying this mapping to a fraction α^* of output. We show that α^* is increasing in the agent's capacity. Thus, expanding the set of feasible distributions by removing a capacity constraint increases the strength of incentives.

The proof goes as follows. Starting at 1, imagine scaling down output by a factor α until the capacity constraint stops binding for at least one optimal contract. Denote this value α^* . For all α in $[\alpha^*, 1]$ we show that the set of optimal contracts is the same and for α^* there is an optimal contract that remains optimal after doing away with the capacity constraint.

Subject to optimal contracts existing, we do not impose a particular structure on the sets of feasible distributions and feasible contracts; the agent's cost function

need only be continuous. This accommodates a wide range of restrictions on allowable contracts (e.g. liability limits, monotonicity, inability to discriminate between collections of states, etc) and a wide range of restrictions on how the agent can influence the distribution over states.

We may apply our result to the literature on capital structure (e.g. [43]; see [44] for a review article). A sector of that literature seeks to derive capital structure from agency between investors and those running firms (e.g. [36]). A common result is that optimality requires the firm to raise capital exclusively through the sale of debts (see [35]). Puzzling, then, that firms sell equity, too. An application of our result resolves this puzzle: the optimal contract under a capacity constraint is debt and a fraction $1 - \alpha^*$ of equity, the remaining equity kept inside the firm to motivate those running it.

2.2 Framework

Output depends on the realization of one among a set of *states* Ω as described by an *output function* $y : \Omega \rightarrow \mathbb{R}$. An agent has the ability to choose the distribution $p \in \Delta(\Omega)$. A principal hires the agent to choose this distribution but cannot verify the choice. The state is verifiable and the agent is given incentives to choose one distribution or another with a *contract*

$$b : \Omega \rightarrow \mathbb{R}.$$

A *profile* (b, p) consists of a contract b and distribution p . The agent's preferences over profiles (b, p) are represented by the expected value of contract payments less a distribution cost:

$$E_p[b] - c(p) = \int b(\omega) dp(\omega) - c(p).$$

The function $c : \Delta(\Omega) \rightarrow \mathbb{R}$ is the agent's *cost function*. The principal's preferences over profiles (b, p) are represented by the expected value of output less contract payments:

$$E_p[y - b] = \int (y(\omega) - b(\omega)) dp(\omega).$$

Contracts are restricted to belong to a set \mathcal{B} of *feasible contracts*. Distributions are restricted to belong to a set \mathcal{D} of *feasible distributions*.

The agent faces a *capacity constraint* which means that included in the restrictions that define the set of feasible distributions \mathcal{D} is the condition that the cost of each is no more than the agent's capacity $k \in \mathbb{R}$:

$$c(p) \leq k.$$

Profile (b, p) is *feasible* if the contract b is feasible, the distribution p is feasible, and this profile is the agent's most preferred profile among those consisting of contract b and any feasible distribution. Profile (b, p) may be *improved upon* if there exists a feasible profile that makes either the agent or principal better off without making the other worse off. *Pareto optimal profiles* are those that are feasible and cannot be improved upon.

2.3 Examples

Many principal-agent setups fit this framework.

The textbook principal-agent problem

In the textbook model, the set of feasible distributions \mathcal{D} is parametrized by effort $e \in \mathbb{R}^n$ and the distribution over output corresponding to effort e is denoted

$p(\cdot|e)$. The agent's cost function maps effort e to its cost $c(p(\cdot|e))$. The set of feasible contracts \mathcal{B} is often defined by liability limits.

Information acquisition

One may be more explicit about how the agent influences the distribution over states. In problems of information acquisition a state ω is taken to be a decision d chosen by the agent and a state of nature θ that the agent cannot influence but can acquire information about. The agent faces a statistical decision problem. He acquires information about the state of nature θ with an *experiment* $(X, \{q(\cdot|\theta)\}_{\theta \in \Theta})$ that consists of a set of signals X and a collection of probability distributions $\{q(\cdot|\theta)\}_{\theta \in \Theta} \subseteq \Delta(X)$ over signals, and he chooses a decision rule $f : X \rightarrow \Delta(D)$, mapping signals to randomizations over decisions. Given a prior π on the states of nature, an experiment and decision rule induce a distribution $p(\cdot|\pi, X, \{q(\cdot|\theta)\}_{\theta \in \Theta}, f) \in \Delta(\Omega)$ over states.

Multiple periods

The distribution over states may arise from the agent's behavior in multiple periods. One could take the states Ω as the set of terminal histories $(\theta_1, \theta_2, \dots, \theta_T)$ where each $\theta_t, t = 1, 2, \dots, T$ is information revealed to the agent at time t . For each history $(\theta_1, \theta_2, \dots, \theta_{t-1})$ the agent chooses a distribution $q(\cdot|\theta_1, \theta_2, \dots, \theta_{t-1})$ over θ_t . This induces a distribution $p(\cdot|q) \in \Delta(\Omega)$.

2.4 Result

Consider a perturbation of our model, parametrized by $\alpha \in [0, 1]$, in which output y in the principal's objective function is replaced by αy (the feasible contracts,

distributions, and profiles stay the same). Denote the set of Pareto optimal profiles to the perturbed problem by $\mathcal{P}(\alpha)$. Further denote by $\mathcal{P}(\alpha, r)$ the subset of $\mathcal{P}(\alpha)$ in which the agent receives the lowest level of utility among the elements of $\mathcal{P}(\alpha)$ that is at least equal to r .

Proposition 4. *Let $(b, p) \in \mathcal{P}(1, r)$ and define $\alpha^* \in [0, 1]$ as*

$$\alpha^* = \sup \{ \alpha \in [0, 1] : c(p') < k \text{ for all } (b', p') \in \mathcal{P}(\alpha, E_p[b] - c(p)) \}.$$

Then, for each $\alpha \in [\alpha^, 1]$,*

$$\mathcal{P}(\alpha, E_p[b] - c(p)) \supseteq \mathcal{P}(1, r),$$

and, conversely, if $(b_\alpha, p_\alpha) \in \mathcal{P}(\alpha, E_p[b] - c(p))$ is such that $c(p_\alpha) = k$, then

$$(b_\alpha, p_\alpha) \in \mathcal{P}(1, r).$$

If the agent's cost function is continuous, then there is a Pareto optimal profile in $\mathcal{P}(\alpha^, E_p[b] - c(p))$ that solves the problem obtained by removing the capacity constraint from the perturbed problem for α^* .*

The point is that all solutions $\mathcal{P}(1, r)$ to the unperturbed problem (a problem in which the capacity constraint might bind) may be obtained as solutions $\mathcal{P}(\alpha^*, E_p[b] - c(p))$ to the perturbed problem for α^* (and for at least one solution of this problem the capacity constraint does not bind).

2.5 Proof

By assumption $(b, p) \in \mathcal{P}(1, r)$. Let $\alpha \in [0, 1)$ (the argument for $\alpha = 1$ is obvious) and $(b_\alpha, p_\alpha) \in \mathcal{P}(\alpha, E_p[b] - c(p))$. We will show:

$$E_p[y] - E_{p_\alpha}[y] \geq E_p[b] - E_{p_\alpha}[b_\alpha] \geq \alpha(E_p[y] - E_{p_\alpha}[y]) \geq 0, \quad (2.1)$$

$$c(p) - c(p_\alpha) \geq E_p[b] - E_{p_\alpha}[b_\alpha]. \quad (2.2)$$

Then, if $\alpha \geq \alpha^*$ we may assume that $c(p_\alpha) = k$, and the inequalities in [A.3](#) give $c(p) = k$ and $E_p[b] = E_{p_\alpha}[b_\alpha]$, and so the inequalities in [A.2](#) give $E_p[y] = E_{p_\alpha}[y]$. Therefore $(b, p) \in \mathcal{P}(\alpha, E_p[b] - c(p))$ (i.e., $\mathcal{P}(\alpha, E_p[b] - c(p)) \supseteq \mathcal{P}(1, r)$) and $(b_\alpha, p_\alpha) \in \mathcal{P}(1, r)$.

To obtain the inequalities in [A.2](#), note $E_p[y - b] \geq E_{p_\alpha}[y - b_\alpha]$ and $E_{p_\alpha}[\alpha y - b_\alpha] \geq E_p[\alpha y - b]$. The first and second inequalities of [A.2](#) are rearrangements of these; and, from this, the last emerges because $\alpha \in [0, 1)$. Inequality [A.3](#) follows from rearrangement of the participation constraint $E_{p_\alpha}[b_\alpha] - c(p_\alpha) \geq E_p[b] - c(p)$.

Finally, we show there is a Pareto optimal profile in $\mathcal{P}(\alpha^*, E_p[b] - c(p))$ that solves the problem obtained by removing the capacity constraint from the perturbed problem for α^* . Let $\alpha_1, \alpha_2, \dots$ be a sequence in $[0, \alpha^*)$ converging to α^* and let $(b_1, p_1), (b_2, p_2), \dots$ be a sequence of profiles with $(b_i, p_i) \in \mathcal{P}(\alpha_i, E_p[b] - c(p))$, $i = 1, 2, \dots$. Since the agent's cost function is continuous the set of feasible profiles yielding the agent at least utility $E_p[b] - c(p)$ is compact. Therefore, there is a subsequence $(b_{n_1}, p_{n_1}), (b_{n_2}, p_{n_2}), \dots$ converging to some (b', p') . Continuity of the principal's objective function implies $(b', p') \in \mathcal{P}(\alpha^*, E_p[b] - c(p))$. Continuity of the agent's cost function together with the fact that $c(p_i) < k$ (because $\alpha_i < \alpha^*$) for $i = 1, 2, \dots$ implies (b', p') solves the problem obtained by removing the capacity constraint from the perturbed problem for α^* . \square

2.6 Application to a theory of capital structure

Suppose that in the absence of a capacity constraint all Pareto optimal contracts correspond to debt contracts

$$y(\omega) - b(\omega) = \min\{y(\omega), F\},$$

for some $F \geq 0$, so that $b(\omega) = \max\{0, y(\omega) - F\}$.

By Proposition 4, if the agent's cost function is continuous, then provided $\mathcal{P}(1, r)$ is nonempty at least one of its elements (b, p) has the form $b(\omega) = \max\{0, \alpha^* y(\omega) - F\} = \alpha^* \max\{0, y(\omega) - F/\alpha^*\}$.

The corresponding payoff for the principal is a debt of F/α^* and a fraction $1 - \alpha^*$ of the remaining equity:

$$y(\omega) - b(\omega) = \min\{y(\omega), F/\alpha^*\} + (1 - \alpha^*) \max\{0, y(\omega) - F/\alpha^*\}.$$

By its definition, α^* increases with agent capacity. Thus, as agent capacity increases, more debt and less equity is issued.

Note that our result does not assume a particular form for the feasible contracts. For example, we could apply our result in a setting where contracts are restricted to be monotonically increasing in output with a slope less and one. Our result also does not assume the form of the feasible set of distributions so we could also apply it if the set distributions were restricted to satisfying a monotone likelihood ratio property. Innes [35] uses both assumptions in deriving debt contracts as optimal and thus we can apply our result to his model.

2.6.1 Live-or-Die Contracts

Dropping contract monotonicity, Innes [35] finds that the resulting optimal contract corresponds to a live-or-die security:

$$y(\omega) - b(\omega) = \begin{cases} y(\omega) & \text{if } y(\omega) < l \\ 0 & \text{if } y(\omega) \geq l \end{cases}.$$

Then

$$b(\omega) = \begin{cases} 0 & \text{if } y(\omega) < l \\ y(\omega) & \text{if } y(\omega) \geq l \end{cases}.$$

The argument applied to the debt and equity case applies here too. In the capacity constrained problem if output is α^* of itself then (i) there is an optimal contract such that the capacity constraint does not bind and (ii) this contract is optimal for in the original problem (the problem with unaltered output). Therefore the agent's contract will have the form shown here but with output replaced by a fraction α^* of itself. The resulting contract form for the principal is

$$y(\omega) - b(\omega) = \begin{cases} y(\omega) & \text{if } y(\omega) < l \\ (1 - \alpha^*)y(\omega) & \text{if } y(\omega) \geq l \end{cases}.$$

Under a capacity constraint the principal retains the fraction $1 - \alpha^*$ of output in the "die" region of the live-or-die contract since he only needs to exhaust the agent's capacity, and this can be done with the fraction α^* of output.

2.7 Discussion

The capital structure irrelevance theorems of Modigliani and Miller [43] state that the market value of a firm is invariant to capital structure because buyers of securities

may borrow or lend making it as if they bought from a firm with any other capital structure.

Many theories now exist to explain why some capital structures are preferred to others (see [44] for a review). These theories relax the assumptions of the Modigliani Miller theorems: the absence of borrowing constraints and taxes, that only debt and equity may be issued, or that the cash flows of the firm are independent from its capital structure.

My results relate to the agency theory of capital structure in which capital structure determines the incentives of the owners of a firm and therefore its cash flows. This point of view was first taken in [36] by Jensen and Meckling who considered a firm's incentives for issuing debt versus equity. Debt is good because it leaves intact entrepreneur incentives whenever firm profits exceed the face value of debt. Equity is bad because it scales down the entrepreneur's incentives: an entrepreneur that sells a 5 percent equity stake will only exert effort to the point where the marginal cost of additional effort is equal to the marginal benefit of an additional 95 cents, not the full 100 cents their additional effort generates. After observing this, Jensen and Meckling ask why firms issue equity at all?

An ingenious entrepreneur eager to expand, has open to him the opportunity to design a whole hierarchy of fixed claims on assets and earnings, with premiums paid for different levels of risk. Why don't we observe large corporations individually owned with a tiny fraction of the capital supplied by the entrepreneur in return for 100 percent of the equity and the rest simply borrowed?

Their answer was that too much debt encourages the entrepreneur to take greater risks in order to obtain the high payoffs where their incentives lie (the alternative being the minimum payoff implied by limited liability). Thus, equity is issued because at some point the harm caused from the risk taking induced by additional

debt is worse than the incentive dilution caused by issuing more equity.

My results suggest an alternative reason why firms issue equity: why leave the entrepreneur with 100 percent equity when a lesser fraction will do? That is, if the entrepreneur is capacity constrained, then selling a certain amount of equity does not change their incentives and so is the best way to initially raise capital. At some point the capacity constraint stops binding. If at this point additional capital is necessary, then it becomes optimal to issue debt.

2.8 Conclusion

We have shown that in solving for optimal incentives imposing a capacity constraint on the agent is equivalent to scaling down output by a factor $\alpha^* \in [0, 1]$ and that this factor is increasing in the agent's capacity. When applied to models that yield debt as the optimal contract a capacity constraint implies that the optimal contract becomes a debt and a fraction $1 - \alpha^*$ of equity. This contributes to the literature that seeks to explain capital structure from a principal-agent relationship between investors and those running the firm.

It would be interesting to find an analogue of our results for when the agent is risk-averse or when the feasible set of actions for the agent changes in a way other than through a capacity constraint.

Chapter 3

Core Equivalence with Large Agents

Competitive equilibrium assumes price taking behavior. A long line of research dating back to Edgeworth [23] has sought to provide a foundation for this assumption by constructing models in which the ability to interact strategically necessarily leads to a competitive outcome (e.g. [19, 5, 25]).

Debreu and Scarf [19] showed that every core allocation is eventually competitive as the size of the economy increases through replication, and Aumann [5] showed that if one models the set of agents as an atomless measure space, then the core and competitive allocations coincide. Subsequent literature (e.g. [50, 24]) has studied settings in which core allocations are necessarily competitive even though the underlying set of agents contains some atoms (i.e., large agents, or groups of agents that act collusively).

We add to this literature by constructing a class of economies in which a large agent exploits the rest of the economy in such a way that the resulting core allocation cannot be decentralized with prices.

3.1 Framework

The economy. Our model of an economy consists of a set of consumers each of whom possess a commodity bundle and hold a preference relation over the set of all commodity bundles. We call the space \mathbb{R}^ℓ the *commodity space*, where ℓ refers to the number of commodities. A *commodity bundle* is an element of the commodity space. A *preference relation*, denoted \succeq , is a linear ordering¹ over the commodity space and induces a *strict preference relation* (denoted \succ) in the following way: If x and y are arbitrary commodity bundles, then $x \succ y$ if and only if $x \succeq y$ and not $y \succeq x$. The set of all such preference relations is \mathcal{P} . The relation $x \succeq y$ is read “ x is preferred or indifferent to y .”

The set of consumers in our economy is T . Let (T, \mathcal{F}, μ) be a measure space. An *allocation* is an integrable function from T into the non-negative orthant of the commodity space. There is a fixed allocation, denoted ω , that is called the *initial allocation*; $\omega(t)$ is called the *initial endowment* of consumer t . An allocation x is said to be *feasible* if $\int x = \int \omega$.

Each consumer, represented by an element of T , has a preference relation defined by a *preference function* $\succeq: T \rightarrow \mathcal{P}$. We say that $\succeq: T \rightarrow \mathcal{P}$ is a *measurable preference function* if, given arbitrary allocations x and y , the set $\{t \in T : x(t) \succ_t y(t)\}$ is measurable.

Definition 1. An economy, denoted \mathcal{E} , is a triad $((T, \mathcal{F}, \mu), \omega, \succeq)$. It consists of a positive and finite measure space (T, \mathcal{F}, μ) , an initial allocation $\omega : T \rightarrow \mathbb{R}_+^\ell$, and a measurable preference function $\succeq: T \rightarrow \mathcal{P}$.

Let (T, \mathcal{F}, μ) be a measure space. An *atom* is defined as a non-null element of the σ -algebra with the property that any subset also belonging to the σ -algebra is of

¹A complete, transitive, and reflexive relation.

equal or zero measure. That is, $A \in \mathcal{T}$ is an atom if for all $B \subseteq A$, where $B \in \mathcal{T}$, we have $\mu(B) = \mu(A)$, or otherwise $\mu(B) = 0$. A measure space with atoms is called *atomic*. A measure space without atoms is called *atomless*.

Core and competitive equilibrium. Fix an economy \mathcal{E} . A *price system*, denoted p , is a vector in \mathbb{R}_+^ℓ .

Definition 2. A *competitive equilibrium* (p, x) is a price system p and an allocation x such that x is feasible, and for μ -almost all t in T

1. $p \cdot x(t) \leq p \cdot \omega(t)$ and
2. $y \succ_t x(t)$ implies $p \cdot y > p \cdot \omega(t)$.

If (p, x) is a competitive equilibrium, then p is called an *equilibrium price system* and x is called a *competitive allocation*. The set of competitive allocations is denoted $W(\mathcal{E})$.

The second equilibrium concept we consider is called the core. A *coalition* is any set S in \mathcal{T} with positive measure. We say that an allocation x is *blocked* by a coalition S if there exists an allocation y such that $\int_S y = \int_S \omega$, and $y(t) \succ_t x(t)$ for all t in S .

Definition 3. The *core* of \mathcal{E} is the set of feasible allocations that are not blocked by any coalition.

The core is denoted $C(\mathcal{E})$.

Non-trivial linear economies. We call an economy *linear* if every consumer's preference relation can be represented by a linear function $u_t : \mathbb{R}_+^\ell \rightarrow \mathbb{R}$ defined by $u_t(x) = a_t \cdot x$, where a_t is a vector in \mathbb{R}_{++}^ℓ such that $\sum_{i=1}^\ell a_{ti} = 1$. An allocation for an economy is called *Pareto optimal* if the coalition consisting of all consumers does

not block the allocation. We call an economy *non-trivial* if the initial allocation is not Pareto optimal.

Our first result provides a motivation for considering non-trivial economies. Recall that we want to characterise economies for which the core does not coincide with the set of competitive allocations. It happens that (under fairly weak conditions on preferences) if the initial endowment is Pareto optimal then the core coincides with the competitive allocations. Thus we exclude these economies from the analysis.

Proposition 5. *Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ be an economy satisfying*

1. *Monotonicity: let x and y be arbitrary commodity bundles. Then $x > y^2$ implies $x \succ_t y$;*
2. *Continuity: let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_t z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_t x\}$ are open;*
3. *$\int_T \omega(t) \gg 0$ (i.e. every commodity exists); and*
4. *\mathcal{E} has at most finitely many atoms and each atom has convex preferences.*

If ω is Pareto optimal, then $C(\mathcal{E}) = W(\mathcal{E})$.

Proof. Consider the allocation $f : T \rightarrow \mathbb{R}_+^\ell$ assigning almost all agents a commodity bundle preferred or indifferent to their initial endowment and a non-null set of agents a commodity bundle preferred to their initial endowment. Let F be the set of integrals of all such allocations. Let I be the set of integrals over the atomless part of the economy of all such allocations. Let G be the set of integrals over the atomic part of the economy of all such allocations. We have that $F = I + G$. By Lyapunov's convexity theorem I is convex. By the convexity of each atom's preferences and the

²Let x and y be arbitrary points in \mathbb{R}^ℓ . We write $x > y$ to mean that $x \neq y$ and $x_i \geq y_i$ for all $i = 1, \dots, \ell$.

assumption that there are finitely many atoms G is convex. Thus F is convex as the sum of two convex sets. F is disjoint from $\int \omega$ because ω is Pareto optimal. By Minkowski's separating hyperplane theorem there exists a price system p such that $x \in F$ implies $p \cdot x \geq p \cdot \int \omega$. We claim (p, ω) is a competitive equilibrium. Suppose not. Then there exists a non-null set of agents S and a commodity bundle $y(t)$ for each agent in S such that $y(t) \succ_t \omega(t)$ and $p \cdot y(t) \leq p \cdot \omega(t)$. Define $z : T \rightarrow \mathbb{R}^\ell$ by

$$z(t) = \begin{cases} \omega(t) & \text{if } t \in T \setminus S, \\ y(t) & \text{if } t \in S. \end{cases}$$

We have $p \cdot \int z \geq p \cdot \int \omega$. If this holds with equality, then for almost all agents in S we have $p \cdot y(t) = p \cdot \omega(t)$. By continuity $y(t) - (\epsilon, \dots, \epsilon) \succ_t \omega(t)$ for sufficiently small $\epsilon > 0$. Define

$$\bar{z}(t) = \begin{cases} \omega(t) & \text{if } t \in T \setminus S, \\ y(t) - (\epsilon, \dots, \epsilon) & \text{if } t \in S. \end{cases}$$

By monotonicity and linearity of the integral we have $p \cdot \int \bar{z} < p \cdot \int \omega$, contradicting that $\int \bar{z}$ belongs to F . Thus

$$p \cdot \int z > p \cdot \int \omega$$

which implies $p \cdot \int_S y > p \cdot \int_S \omega$, contradicting $p \cdot y(t) \leq p \cdot \omega(t)$ for all $t \in S$. Therefore (p, ω) is a competitive equilibrium.

We now apply the argument of [5] to show that every core allocation is a competitive allocation. For any feasible allocation x let Γ_t^x be the set of all commodity bundles z in the commodity space such that $z + \omega(t) \succ_t x(t)$ and let Γ^x be the convex hull of the union of each Γ_t^x , where the union is taken over any measurable set U such that $\mu(T \setminus U) = 0$. As in [5] x is a competitive allocation if and only if the origin does not belong to Γ^x . Now, let x be a core allocation. Since $x(t) \succeq_t \omega(t)$

μ -a.e. we have $\Gamma_t^x \subseteq \Gamma_t^\omega$ and thus $\Gamma^x \subseteq \Gamma^\omega$. Since ω is a competitive allocation the origin does not belong to Γ^ω , and thus also not to Γ^x . Hence x is a competitive allocation. \square

Unbalanced economies. Two consumers are said to be of the same *type* if they have the same preference relation and the same initial endowment. An economy is said to be an *n-type* economy if the set of consumers can be partitioned into n non-null sets, each containing exactly those consumers of a particular type. The economy is called *unbalanced* if an element of this partition consists of a single atom, i.e. there is an atom whose type is not shared by any other non-null set of agents.

3.2 Analysis and results

Our main result is:

Theorem. *Let (T, \mathcal{T}, μ) be a measure space and a_1 and a_2 distinct vectors in \mathbb{R}_{++}^ℓ such that $\sum_{i=1}^\ell a_{ti} = 1$. Consider the set of two-type non-trivial linear economies on the measure space (T, \mathcal{T}, μ) such that the preference relation of type i consumers is represented by $u_i(x) = a_i \cdot x$. If for all such economies \mathcal{E} we have $C(\mathcal{E}) \subseteq W(\mathcal{E})$, then (T, \mathcal{T}, μ) is atomless.*

Equivalently:

Theorem. *Let \mathcal{E} be a two-type unbalanced non-trivial linear economy. There exists a non-competitive core allocation for \mathcal{E} .*

The basic idea of the proof can be summarised as follows. Let \mathcal{E} be a two-type unbalanced non-trivial linear economy. We first show that if one restricts attention to a certain class of allocations, then \mathcal{E} is equivalent to a non-trivial linear economy

with two atoms. We then show that for any non-trivial linear economy with two atoms there exists a core allocation that is not competitive. Finally we show that this allocation is also a non-competitive core allocation for \mathcal{E} .

Interpretation of the main theorem. Our main result is that for every two-type unbalanced non-trivial linear economy the core is a strict superset of the set of competitive allocations. There are two ways to interpret this result. The first is mathematical and the second economic.

The mathematical interpretation is that our result provides a necessary and sufficient condition for a positive and finite measure space to be atomless. Suppose we are given a positive finite measure space and distinct vectors a_1 and a_2 in \mathbb{R}_{++}^ℓ such that $\sum_{i=1}^\ell a_{ti} = 1$. Consider the set of two-type non-trivial linear economies on this measure space such that the preference relation of type i consumers is represented by $u_i(x) = a_i \cdot x$. A necessary and sufficient condition for this measure space to be atomless is that all such economies have the property that the core is a subset of the set of competitive allocations.

The economic interpretation is as follows: Our result shows that given any positive and finite measure space there exists economies on the measure space such that core equivalence does not hold. The presence of a large agent of a unique type in the economy can cause the core equivalence result to break down.

The important condition of the class of economies we identify is that they are unbalanced: that there is an atom whose type is not shared by any other consumers in the economy. In [50] core equivalence is obtained only if there are at least two atoms of the same type. In [24] each atom must be associated with a non-null set of atomless consumers of the same type. These results require that the economy not be unbalanced. In an economy with two identical firms we are likely to observe

intense competition between them, but in an economy with a firm of a unique type we will observe monopolistic behaviour.

3.2.1 A rescaling lemma

This section presents a result of [42] which we will use to simplify the analysis of following sections. The basic idea is that we can transform one economy to another in such a way that the relationship between the core and competitive equilibria remains fixed. The transformation is attractive because it gives us the freedom to choose the measure of a finite number of measurable subsets of consumers. When working with a particular economy we will find it useful to transform the economy so that certain measurable subsets of consumers are of measure one.

The transformation is also interesting in its own right. It implies that any economy defines a class of ‘equivalent’ economies, from the point of view of core equivalence analysis. That is, any property of the core and competitive equilibria that holds in one economy from the class must hold in all economies of that class.

The transform. Fix a set of consumers T . Let \mathbb{E} denote the class of all economies in which the set of consumers is T and Λ the class of all simple functions $\lambda = \sum_{i=1}^n \lambda_i \chi_{G_i}$ from the set of consumers to the positive real numbers. Define a mapping $\Xi : \mathbb{E} \times \Lambda \rightarrow \mathbb{E}$ as follows. For an economy $\mathcal{E} = ((T, \mathcal{F}, \mu), \omega, \succeq)$ in \mathbb{E} and a simple function λ in Λ let $\Xi(\mathcal{E}, \lambda)$ be the economy $((T, \mathcal{F}, \hat{\mu}), \hat{\omega}, \hat{\succeq})$ in \mathbb{E} defined by:

1. $\hat{\omega} = \lambda\omega$;
2. $\hat{\succeq}_t$ is defined by $\lambda(t)x \hat{\succeq}_t \lambda(t)y$ if and only if $x \succeq_t y$; and

3. $\hat{\mu}$ is defined by the equation

$$\hat{\mu}(S) = \sum_{i=1}^n \frac{\mu(S \cap G_i)}{\lambda_i}$$

for any S in \mathcal{T} .

Notice that $\hat{\mu}$ is a measure on (T, \mathcal{T}) . Also, it is clear that if x is an allocation and λ an element of Λ , then λx is integrable. Thus the mapping is well defined. We now prove the main result of this section: that the class of economies $\Xi(\mathcal{E}, \Lambda)$ are equivalent in the following sense.

Lemma 1. *Let \mathcal{E} be an economy. For any λ in Λ we have:*

1. $x \in C(\mathcal{E})$ if and only if $\lambda x \in C(\Xi(\mathcal{E}, \lambda))$, and
2. $x \in W(\mathcal{E})$ if and only if $\lambda x \in W(\Xi(\mathcal{E}, \lambda))$.

Proof. First it is shown that an allocation x for \mathcal{E} is feasible if and only if λx is feasible for $\Xi(\mathcal{E}, \lambda)$. To see this let x be an allocation for \mathcal{E} . Then

$$\int \lambda x d\hat{\mu} = \int \lambda x \frac{d\mu}{\lambda} = \int x d\mu.$$

To see that (1) is true let x be a feasible allocation for \mathcal{E} . Suppose x can be blocked by a coalition S . Then there exists an allocation y for \mathcal{E} such that $\int_S y = \int_S \omega$ and $y(t) \succ_t x(t)$ for all t in S . By the previous observation λy is a feasible allocation for $\Xi(\mathcal{E}, \lambda)$ and $\int_S \lambda y d\hat{\mu} = \int_S \lambda \omega d\hat{\mu}$. By the definition of $\hat{\mu}$ we have $\lambda(t)y(t) \succ_t \lambda(t)x(t)$ for all t in S . Thus λx can be blocked by S in $\Xi(\mathcal{E}, \lambda)$. The converse holds by symmetry.

To see that (2) is true let (p, x) be a competitive equilibrium for \mathcal{E} . Then for μ -almost all consumers t in T we have $p \cdot x(t) \leq p \cdot \omega(t)$, and $y \succ_t x(t)$ implying $p \cdot y > p \cdot \omega(t)$. Since $\lambda(t)$ is a positive real number $p \cdot \lambda(t)x(t) \leq p \cdot \lambda(t)\omega(t)$ μ -a.e.

By the definition of \succeq it is clear that $\lambda(t)y \succeq_t \lambda(t)x(t)$ implies $p \cdot \lambda(t)y > p \cdot \lambda(t)\omega(t)$ μ -a.e. Thus $(p, \lambda x)$ is a competitive equilibrium for $\Xi(\mathcal{E}, \lambda)$. Once again the converse holds by symmetry. \square

3.2.2 A two-atom economy

This section studies the class of non-trivial linear economies with two atoms. By Lemma 1 we can assume without loss of generality that the measure of each atom is one. Consequently the analysis devolves to a simple two consumer economy of the type studied by Edgeworth. The main result of this section says that if the price system facing each atom is the vector defining the other atom's utility function, then the solution to one of the resulting utility maximisation problems is infeasible with the resources of the economy.

Preliminary observations. Let \mathcal{E} be a non-trivial linear economy with two atoms (note that such an economy is necessarily two-type and unbalanced.) By Lemma 1 we may assume without loss of generality that the measure of each atom is one. Denote the atoms A_1 and A_2 . Since the economy is linear each atom's preference relation can be represented by a linear utility function $u_i : \mathbb{R}_+^\ell \rightarrow \mathbb{R}$ defined by $u_i(x) = a_i \cdot x$ where a_i belongs to \mathbb{R}_+^ℓ for both $i = 1$ and $i = 2$ such that $\sum_{i=1}^2 a_{ti} = 1$. Since the economy is non-trivial $a_i \in \mathbb{R}_{++}^\ell$ and $a_1 \neq a_2$ (for otherwise, ω would be Pareto optimal.) Fix a price system p and consider the following utility maximisation problem.

Problem 1.

$$\max_{x \in \beta_i(p)} a_i \cdot x, \text{ where}$$

$$\beta_i(p) = \{x \in \mathbb{R}_+^\ell : p \cdot x \leq p \cdot \omega(A_i)\}$$

Let x be a commodity bundle. Define the *support* of x to be

$$\text{supp } x = \{l : x_l > 0\}.$$

For each A_i , $i = 1, 2$, define the set of commodities yielding maximum marginal utility by $S(p, a_i)$. Thus

$$S(p, a_i) = \{l : a_l p_k \geq a_k p_l \text{ for all } k = 1, \dots, \ell\}.$$

If x^* is a solution to Problem 1, then it is clear that each atom consumes only those commodities yielding maximum marginal utility. That is,

Lemma 2. *If x^* is a solution to Problem 1, then $\text{supp } x^* \subseteq S(p, a_i)$ for $i = 1$ and $i = 2$.*

We now prove that if each atom A_i faces prices $p = a_j$ with $j \neq i$, then the set of commodities that yield maximum marginal utility for each are disjoint.

Lemma 3. $S(a_2, a_1) \cap S(a_1, a_2) = \emptyset$.

Proof. Suppose not. Then there exists a commodity l belonging to both $S(a_2, a_1)$ and $S(a_1, a_2)$ so that $a_{1l}a_{2k} \geq a_{1k}a_{2l}$ and $a_{2l}a_{1k} \geq a_{2k}a_{1l}$ for all $k = 1, \dots, \ell$. This implies

$$a_{1l}a_{2k} = a_{1k}a_{2l} \text{ for all } k = 1, \dots, \ell.$$

If $a_{1l} = a_{2l}$, then $a_1 = a_2$, a contradiction. Thus there exists a positive number α , not equal to one, such that $a_{1l} = \alpha a_{2l}$. But then $a_{1k} = \alpha a_{2k}$ for all $k = 1, \dots, \ell$ so that

$$1 = \sum_{k=1}^{\ell} a_{1k} = \alpha \sum_{k=1}^{\ell} a_{2k} = \alpha.$$

A contradiction. □

Decentralization. The main result of this section is that if atoms A_1 and A_2 each face a price system defined by the other atom's utility function, then the solution to

Problem 1 is infeasible with the total resources of the economy for at least one atom.

Lemma 4. *Suppose x_1^* solves Problem 1 for A_1 with $p = a_2$ and suppose that x_2^* solves Problem 1 for A_2 with $p = a_1$. Then, for at least one $i = 1, 2$ we have $x_i^* \not\leq \omega(A_1) + \omega(A_2)$.*

Proof. Suppose by contradiction that for both $i = 1$ and $i = 2$ we have $x_i^* \leq \omega(A_1) + \omega(A_2)$. By Lemma 2 we have $\text{supp } x_1^* \subseteq S(a_2, a_1)$ and $\text{supp } x_2^* \subseteq S(a_1, a_2)$. By Lemma 3 $S(a_2, a_1) \cap S(a_1, a_2) = \emptyset$. Thus

$$\text{supp } x_1^* \cap \text{supp } x_2^* = \emptyset.$$

Since $x_i^* \leq \omega(A_1) + \omega(A_2)$ for both $i = 1$ and $i = 2$ we have

$$x_1^* + x_2^* \leq \omega(A_1) + \omega(A_2).$$

By Walras' law $a_1 \cdot x_2^* = a_1 \cdot \omega(A_2)$. Since ω is not Pareto optimal there exists commodity bundles x_1 and x_2 such that $x_1 + x_2 = \omega(A_1) + \omega(A_2)$, $a_1 \cdot x_1 > a_1 \cdot \omega(A_1)$ and $a_2 \cdot x_2 > a_2 \cdot \omega(A_2)$. x_1 is affordable since $a_2 \cdot (x_1 + x_2) = a_2 \cdot (\omega(A_1) + \omega(A_2))$ implies

$$a_2 \cdot x_1 < a_2 \cdot \omega(A_1).$$

Thus $a_1 \cdot x_1^* > a_1 \cdot \omega(A_1)$. Combining these two observations yields

$$a_1 \cdot (x_1^* + x_2^*) > a_1 \cdot (\omega(A_1) + \omega(A_2)),$$

which contradicts $x_1^* + x_2^* \leq \omega(A_1) + \omega(A_2)$. □

3.2.3 Constructing the core allocation

This section studies the class of two-type unbalanced non-trivial linear economies. The main result of this section is a characterisation of a core allocation for such

economies.

To begin, fix a two-type unbalanced non-trivial linear economy \mathcal{E} . Since \mathcal{E} is unbalanced there is an atom of a unique type. Denote this atom A_1 . We will call consumers belonging to A_1 type one and consumers belonging to $T \setminus A_1$ type two. To simplify notation we denote the initial endowment of a consumer of type i by ω_i . Denote the utility function representing the preferences of a consumer of type i by $u_i(x) = a_i \cdot x$.

It happens that in such an economy there is a core allocation in which all consumers of type two get a commodity bundle indifferent to their initial commodity bundle and the atom of type one maximises his utility over the set of feasible allocations satisfying this constraint. We now construct this allocation. Consider the following maximisation problem.

Problem 2.

$$\begin{aligned} \max_{(x_1, x_2) \in X} u_1(x_1) \\ X = F \cap W \end{aligned}$$

where

$$F = \{(x_1, x_2) \in \mathbb{R}_+^{2\ell} : \mu(A_1)x_1 + \mu(T \setminus A_1)x_2 \leq \mu(A_1)\omega_1 + \mu(T \setminus A_1)\omega_2\}$$

and

$$W = \{(x_1, x_2) \in \mathbb{R}_+^{2\ell} : u_2(x_2) = u_2(\omega_2)\}.$$

The main result of this section is the following theorem.

Lemma 5. *Let \mathcal{E} be a two-type unbalanced non-trivial linear economy. A solution (x_1^*, x_2^*) to Problem 2 exists. The allocation $x^* : T \rightarrow \mathbb{R}_+^\ell$ that assigns x_i^* to each consumer of type i belongs to $C(\mathcal{E})$.*

Proof. A solution exists. Notice that X is closed and bounded. By the Heine-Borel property X is compact. As u_1 is continuous we obtain a solution by Weierstrass' theorem.

Now, let (x_1^*, x_2^*) be a solution to Problem 2 and define x^* as in the statement of the lemma. Clearly x^* is a feasible allocation. Suppose by way of contradiction that x^* does not belong to $C(\mathcal{E})$. Then there exists a coalition S which blocks x^* with an allocation y . Note that A_1 belongs to S . For if not, then by the linearity of u_2 we have that $\frac{1}{\mu(S)} \int_S y$ is preferred to x_2^* by each consumer of type two. However, $\frac{1}{\mu(S)} \int_S y = \omega_2$ contradicting $u_2(x_2^*) = u_2(\omega_2)$. By the same argument a non-null set of consumers of type two belongs to S . Define

$$\bar{y} = \frac{1}{\mu(T \setminus A_1)} \left[\int_{S \setminus A_1} y + \int_{T \setminus S} \omega_2 \right].$$

Clearly \bar{y} is preferred to ω_2 by each consumer of type two. Since u_2 is continuous, by the Intermediate Value Theorem there exists an $\alpha \in (0, 1)$ such that $\alpha \bar{y}$ is indifferent to ω_2 for each consumer of type two. Finally

$$\begin{aligned} \mu(A_1)y(A_1) + \mu(T \setminus A_1)\alpha\bar{y} &\leq \mu(A_1)y(A_1) + \mu(T \setminus A_1)\bar{y} \\ &= \mu(A_1)y(A_1) + \int_{S \setminus A_1} y + \int_{T \setminus S} \omega_2 \\ &= \int \omega \\ &= \mu(A_1)\omega_1 + \mu(T \setminus A_1)\omega_2. \end{aligned}$$

Thus $(y(A_1), \alpha\bar{y})$ belongs to X contradicting that (x_1^*, x_2^*) is a maximiser. \square

3.2.4 Proof of the main theorem

In any two-type unbalanced non-trivial linear economy Lemma 5 characterises a core allocation. It turns out via Lemma 4 that this allocation is not competitive. Thus, in

every two-type unbalanced non-trivial linear economy there exists a non-competitive core allocation. This is the main result.

Theorem 1. *Let \mathcal{E} be a two-type unbalanced non-trivial linear economy. There exists a non-competitive core allocation for \mathcal{E} .*

Proof. By Theorem 1 there is no loss of generality in assuming that $\mu(A_1) = \mu(T \setminus A_1) = 1$. Since \mathcal{E} is a two-type unbalanced economy Theorem 5 yields the core allocation x^* . Suppose by contradiction that x^* is a competitive allocation. We claim that $p = a_2$. For if $p \neq a_2$, then since ω_2 is not Pareto optimal there exists a commodity bundle y such that $p \cdot y \leq p \cdot \omega_2$ and $a_2 \cdot y > a_2 \cdot \omega_2$. Since $a_2 \cdot x_2^* = a_2 \cdot \omega_2$ this contradicts our assumption that x^* is a competitive allocation. Thus, $p = a_2$. By Theorem 4, interchanging a_1 and a_2 if necessary, we have $x_2^* \not\leq \omega_1 + \omega_2$. Thus $\int x^* \not\leq \int \omega$. A contradiction. \square

Finally:

Theorem. *Let (T, \mathcal{F}, μ) be a measure space and a_1 and a_2 distinct vectors in \mathbb{R}_{++}^ℓ such that $\sum_{i=1}^\ell a_{ti} = 1$. Consider the set of two-type non-trivial linear economies on the measure space (T, \mathcal{F}, μ) such that the preference relation of type i consumers is represented by $u_i(x) = a_i \cdot x$. If for all such economies \mathcal{E} we have $C(\mathcal{E}) \subseteq W(\mathcal{E})$, then (T, \mathcal{F}, μ) is atomless.*

Proof. By Theorem 1 if (T, \mathcal{F}, μ) is atomic, then $C(\mathcal{E}) \not\subseteq W(\mathcal{E})$ for some such two-type non-trivial linear economy. \square

3.3 Graphical proof for the two commodity case

After normalizing the economy so that $\mu(A) = \mu(T \setminus A) = 1$, we may represent the two commodity situation in an Edgeworth box with the caveat that for type 2 what

is measured is the per-capita allocation $\int_{T \setminus A} x$. In the diagrams below, the core is in red (hollow dot) and the set of competitive allocations is in blue (solid dot) (not to be confused with the initial allocation).




Figure 3.1: Proof of the main theorem: configuration 1.






Figure 3.2: Proof of the main theorem: configuration 2.

In the following case, we interchange a_1 and a_2



(a)



(b)

Figure 3.3: Proof of the main theorem: configuration 3. Interchange a_1 and a_2 .

References

- [1] George A Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, pages 488–500, 1970. 23
- [2] Kenneth J Arrow. Information and economic behavior. Technical report, DTIC Document, 1973. 24
- [3] Kenneth J Arrow. Informational structure of the firm. *The American Economic Review*, 75(2):303–307, 1985. 24
- [4] Kenneth J Arrow. Agency and the market. *Handbook of mathematical economics*, 3:1183–1195, 1986. 25
- [5] Robert J Aumann. Markets with a continuum of traders. *Econometrica: Journal of the Econometric Society*, pages 39–50, 1964. 36, 40, 91, 102, 116, 122, 123
- [6] Robert J Aumann and Michael Maschler. *Repeated games with incomplete information*. MIT press, 1995. 72
- [7] Dirk Bergemann and Ulrich Hege. Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance*, 22(6):703–735, 1998. 23
- [8] David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, 24(2):265–272, 1953. 7
- [9] Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999. 24
- [10] Andrew Caplin and Mark Dean. Behavioral implications of rational inattention with Shannon entropy. Technical report, National Bureau of Economic Research, 2013. 9, 24, 72, 74, 77
- [11] Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *The American Economic Review*, 105(7):2183–2203, 2015. 24
- [12] Gabriel Carroll. Robustness and linear contracts. *The American Economic Review*, 105(2):536–563, 2015. 2

- [13] Gabriel Carroll. Robust incentives for information acquisition. *working paper*, 2016. [5](#), [21](#)
- [14] Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, 2013. [24](#)
- [15] Jacques Crémer and Fahad Khalil. Gathering information before signing a contract. *The American Economic Review*, pages 566–578, 1992. [23](#)
- [16] Jacques Cremer, Fahad Khalil, and Jean-Charles Rochet. Strategic information gathering before a contract is offered. *Journal of Economic Theory*, 81(1):163–200, 1998. [23](#)
- [17] Tri Vi Dang, Gary Gorton, and Bengt Holmström. Ignorance, debt and financial crises. *Yale University and Massachusetts Institute of Technology, working paper*, 2012. [23](#)
- [18] Gerard Debreu. *Theory of value: An axiomatic analysis of economic equilibrium*. Number 17. Yale University Press, 1959. [101](#), [111](#)
- [19] Gerard Debreu and Herbert Scarf. A limit theorem on the core of an economy. *International Economic Review*, 4(3):235–246, 1963. [36](#), [91](#), [95](#), [108](#), [123](#), [125](#)
- [20] Morris H DeGroot. Uncertainty, information, and sequential experiments. *The Annals of Mathematical Statistics*, pages 404–419, 1962. [24](#)
- [21] Joel S Demski and David EM Sappington. Delegated expertise. *Journal of Accounting Research*, pages 68–89, 1987. [4](#)
- [22] Peter Diamond. Managerial incentives: on the near linearity of optimal compensation. *Journal of Political Economy*, 106(5):931–957, 1998. [2](#)
- [23] Francis Ysidro Edgeworth. *Mathematical psychics: An essay on the application of mathematics to the moral sciences*, volume 10. Kegan Paul, 1881. [36](#)
- [24] Jean Jaskold Gabszewicz and Jean-François Mertens. An equivalence theorem for the core of an economy whose atoms are not "too" big. *Econometrica: Journal of the Econometric Society*, pages 713–721, 1971. [36](#), [42](#), [112](#), [122](#)
- [25] Douglas Gale. Bargaining and competition part i: characterization. *Econometrica: Journal of the Econometric Society*, pages 785–806, 1986. [36](#)
- [26] Dino Gerardi and Lucas Maestri. A principal–agent model of sequential testing. *Theoretical Economics*, 7(3):425–463, 2012. [23](#)
- [27] Joseph Greenberg and Benyamin Shitovitz. A simple proof of the equivalence theorem for oligopolistic mixed markets. *Journal of mathematical economics*, 15(2):79–83, 1986. [117](#)

- [28] Yingni Guo. Dynamic delegation of experimentation. *The American Economic Review*, 106(8):1969–2008, 2016. 23
- [29] Marina Halac, Navin Kartik, and Qingmin Liu. Optimal contracts for experimentation. *The Review of Economic Studies*, 83(3):1040–1091, 2016. 23
- [30] Paul R Halmos. *Measure theory*, volume 18. Springer, 2013. 102
- [31] Benjamin Hébert and Michael Woodford. Rational inattention with sequential information sampling. *working paper*, 2017. 9, 15, 24, 67, 72
- [32] Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *The American economic review*, 61(4):561–574, 1971. 23
- [33] Bengt Holmström and Paul Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, pages 303–328, 1987. 2
- [34] Johannes Hörner and Larry Samuelson. Incentives for experimenting agents. *The RAND Journal of Economics*, 44(4):632–663, 2013. 23
- [35] Robert D Innes. Limited liability and incentive contracting with ex-ante action choices. *Journal of economic theory*, 52(1):45–67, 1990. 3, 27, 32, 33
- [36] Michael C Jensen and William H Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of financial economics*, 3(4):305–360, 1976. 3, 27, 34
- [37] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011. 72, 74
- [38] Godfrey Keller, Sven Rady, and Martin Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, 2005. 24
- [39] Dennis V Lindley. On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005, 1956. 24
- [40] Gustavo Manso. Motivating innovation. *The Journal of Finance*, 66(5):1823–1860, 2011. 23
- [41] Filip Matejka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *The American Economic Review*, 105(1):272–298, 2014. 77
- [42] Jean-François Mertens and Sylvain Sorin. *Game-theoretic methods in general equilibrium analysis*, volume 77. Springer Science & Business Media, 1994. 43, 127

- [43] Franco Modigliani and Merton H Miller. The cost of capital, corporation finance and the theory of investment. *The American economic review*, pages 261–297, 1958. 27, 33
- [44] Stewart C Myers. Capital structure. *The journal of economic perspectives*, 15(2):81–102, 2001. 27, 34
- [45] Henrique Oliveira, Tommaso Denti, Maximilian Mihm, and Kemal Ozbek. Rationally inattentive preferences and hidden information costs. *Theoretical Economics*, 12(2):621–654, 2017. 24
- [46] Kent Osband. Optimal forecasting incentives. *Journal of Political Economy*, 97(5):1091–1112, 1989. 4
- [47] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer, 2009. 65
- [48] Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. In *Foundations of Insurance Economics*, pages 355–375. Springer, 1976. 23
- [49] Leonard J Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971. 4
- [50] Benyamin Shitovitz. Oligopoly in markets with a continuum of traders. *Econometrica: Journal of the Econometric Society*, pages 467–501, 1973. 36, 42, 112, 116
- [51] Christopher A Sims. Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690, 2003. 24
- [52] Christopher A Sims et al. *Rational inattention: a research agenda*. Deutsche Bundesbank, 2005. 24
- [53] Michael Spence. Job market signaling. *The quarterly journal of Economics*, 87(3):355–374, 1973. 23
- [54] Robert M Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic theory*, 21(2):265–293, 1979. 23
- [55] Karl Vind. A third remark on the core of an atomless economy. *Econometrica*, 40(3):585, 1972. 119, 121
- [56] Michael Woodford. Inattentive valuation and reference-dependent choice. *Unpublished Manuscript, Columbia University*, 2012. 24
- [57] Ming Yang. Optimality of debt under flexible information acquisition. *Available at SSRN 2103971*, 2015. 23

- [58] Ming Yang and Yao Zeng. Financing entrepreneurial production: Security design with flexible information acquisition. *Available at SSRN 2194194*, 2015. 23
- [59] Luis Zermeno. A principal-expert model and the value of menus. *unpublished paper, Massachusetts Institute of Technology*, 2011. 5
- [60] Luis Zermeno. The role of authority in a general principal-expert model. *unpublished paper, Massachusetts Institute of Technology*, 2012. 5

Appendix A

Appendix to Chapter 1

A.1 Proofs

A.1.1 Proof of Proposition 1

Let p be an experiment that is optimal for the agent against contract y . First, consider a contract $\alpha y - \beta$ for $\alpha \geq \alpha'$ and $\beta : \Theta \rightarrow \mathbb{R}$ and let p_α be optimal for this contract. We will show that

$$E_p[y] - E_{p_\alpha}[y] \geq c(p) - c(p_\alpha) \geq \alpha(E_p[y] - E_{p_\alpha}[y]). \quad (\text{A.1})$$

Then, if $\alpha \geq \alpha'$, we may assume $c(p_\alpha) = k$, so that $c(p) = k$ and $E_p[y] = E_{p_\alpha}[y]$. This implies that $\alpha y - \beta$ is first-best.

Inequality (A.1) may be proved by noting that $E_p[y] - c(p) \geq E_{p_\alpha}[y] - c(p_\alpha)$ and $E_{p_\alpha}[\alpha y - \beta] - c(p_\alpha) \geq E_p[\alpha y - \beta] - c(p)$. Rearranging these inequalities gives the first and second inequalities of (A.1).

Conversely, let b be a first-best contract and p' an experiment that is optimal for this contract. Then the gradient of the agent's objective function at p' under contract y and contract b belong to the normal cone $N_{\mathcal{E}}(p')$ (see Theorem 2).

Since the cost function is differentiable, the set of feasible experiments is regular at p' . Hence,

$$N_{\mathcal{E}}(p') = \{\mu \nabla c(p') + \rho : \mu \geq 0 \text{ and } \rho : \Theta \rightarrow \mathbb{R}\}.$$

Therefore, there are $\mu, \mu' \geq 0$ and $\rho, \rho' : \Theta \rightarrow \mathbb{R}$ such that

$$\pi y - \nabla c(p') = \mu \nabla c(p') + \rho \text{ and } \pi b - \nabla c(p') = \mu' \nabla c(p') + \rho'.$$

Combining these equations gives

$$b = \frac{1 + \mu'}{1 + \mu} y - \frac{(\mu' + 1)\rho - (\mu + 1)\rho'}{\pi(\mu + 1)}.$$

To complete the argument we need to show that the coefficient on b is at least equal to α' . Suppose, on the contrary, that this coefficient is less than α' . Then, from the definition of α' , $c(p') < k$. But this implies $\mu = \mu' = 0$ so that $(1 + \mu')/(1 + \mu) = 1$, and this is impossible because, by its definition, $\alpha' \leq 1$. \square

A.1.2 Proof of Proposition 2

By assumption $(b, p) \in \mathcal{P}(1, r)$. Let $\alpha \in [0, 1)$ (the argument for $\alpha = 1$ is trivial) and $(b_\alpha, p_\alpha) \in \mathcal{P}(\alpha, E_p[b] - c(p))$. We will show:

$$E_p[y] - E_{p_\alpha}[y] \geq E_p[b] - E_{p_\alpha}[b_\alpha] \geq \alpha(E_p[y] - E_{p_\alpha}[y]) \geq 0, \quad (\text{A.2})$$

$$c(p) - c(p_\alpha) \geq E_p[b] - E_{p_\alpha}[b_\alpha]. \quad (\text{A.3})$$

Then, if $\alpha \geq \alpha^*$ we may assume that $c(p_\alpha) = k$, and the inequalities in (A.3) give $c(p) = k$ and $E_p[b] = E_{p_\alpha}[b_\alpha]$, and so the inequalities in (A.2) give $E_p[y] = E_{p_\alpha}[y]$. Therefore $(b, p) \in \mathcal{P}(\alpha, E_p[b] - c(p))$ (so that $\mathcal{P}(\alpha, E_p[b] - c(p)) \supseteq \mathcal{P}(1, r)$) and $(b_\alpha, p_\alpha) \in \mathcal{P}(1, r)$.

To obtain the inequalities in (A.2) note that $E_p[y - b] \geq E_{p_\alpha}[y - b_\alpha]$ and $E_{p_\alpha}[\alpha y - b_\alpha] \geq E_p[\alpha y - b]$. The first and second inequalities of (A.2) are rearrangements of these; and from this the last inequality in (A.2) follows since $\alpha \in [0, 1]$. Inequality (A.3) follows from the participation constraint $E_{p_\alpha}[b_\alpha] - c(p_\alpha) \geq E_p[b] - c(p)$.

What remains to be shown is that there is a Pareto optimal profile in $\mathcal{P}(\alpha^*, E_p[b] - c(p))$ that solves the problem obtained by removing the capacity constraint from the perturbed problem for α^* . Let $\alpha_1, \alpha_2, \dots$ be a sequence in $[0, \alpha^*)$ converging to α^* and let $(b_1, p_1), (b_2, p_2), \dots$ be a sequence of profiles with $(b_i, p_i) \in \mathcal{P}(\alpha_i, E_p[b] - c(p))$, for $i = 1, 2, \dots$. Since the agent's cost function is continuous, the set of feasible profiles yielding the agent at least utility $E_p[b] - c(p)$ is compact. Therefore, there is a sub-sequence $(b_{n_1}, p_{n_1}), (b_{n_2}, p_{n_2}), \dots$ converging to some feasible (b', p') . Continuity of the principal's objective function implies $(b', p') \in \mathcal{P}(\alpha^*, E_p[b] - c(p))$. Continuity of the agent's cost function together with the fact that $c(p_i) < k$ for $i = 1, 2, \dots$ (because $\alpha_i < \alpha^*$) implies that (b', p') solves the problem obtained by removing the capacity constraint from the perturbed problem for α^* . \square

A.1.3 Proof of Proposition 3

Step 1: necessary and sufficient conditions for optimal experiments

The agent's problem is convex and so a necessary and sufficient condition for an experiment p to be optimal against contract b is that $\nabla_p (E_p[b] - c(p))$ belong to the normal cone $N_{\mathcal{E}}(p)$ (see Theorem 2):

$$\pi(\theta)b(d, \theta) - \frac{\partial c(p)}{\partial p(d|\theta)} = \rho[\theta] \text{ for all } d \in D, \theta \in \Theta \quad (\text{A.4})$$

for some $\rho \in \mathbb{R}^{|\Theta|}$. Here we have used that the set of feasible experiments is defined without a capacity constraint and that the experiment p , by assumption, places

positive probability on each decision in each state.

Step 2: a necessary condition for optimal contracts

Condition [A.4](#) involves Lagrange multipliers ρ so we shall pose the problem of designing an optimal contract as a choice of a contract b , experiment p , and Lagrange multipliers ρ such that the contract and experiment form a feasible profile that yields the agent at least utility r . In the notation of [Theorem 3](#) define $C = \{(b, p, \rho) \in A : F(b, p, \rho) \in B\}$, where $A = \mathbb{R}^{|D \times \Theta| + |D \times \Theta| + |\Theta|}$, $F : \mathbb{R}^{|D \times \Theta| + |D \times \Theta| + |\Theta|} \rightarrow \mathbb{R}^{|D \times \Theta| + 1 + |\Theta| + |D \times \Theta|}$ is defined componentwise $F(b, p, \rho) = (f_j^i(b, p, \rho))_{i,j}$ by

$$f_j^i(b, p, \rho) = \begin{cases} b(d, \theta) & \text{for } i = 1, j \in D \times \Theta, \\ r - (E_p[b] - c(p)) & \text{for } i = 2, \\ \sum_{d \in D} p(d|\theta) - 1 & \text{for } i = 3, j \in \Theta, \\ \pi(\theta)b(d, \theta) - \frac{\partial c(p)}{\partial p(d|\theta)} - \rho[\theta] & \text{for } i = 4, j \in D \times \Theta, \end{cases}$$

and $B = \left(\prod_{(d, \theta) \in D \times \Theta} [0, y(d, \theta)] \right) \times (-\infty, 0] \times \{0\}^{|\Theta| + |D \times \Theta|}$.

The Pareto optimal contracts in which the principal maximizes his utility subject to the agent receiving at least utility r is given by maximizing $E_p[y - b]$ over C . A necessary condition for the local optimality of (b, p, ρ) is that $\nabla_{(b, p, \rho)} E_p[\alpha^* y - b]$ belongs to the normal cone $\hat{N}_C(b, p, \rho)$

The main challenge of the proof is in showing that C is regular at any solution (b, p, ρ) so that we may then express its normal cone as certain linear combinations of the gradients of the constraint functions.

We follow the notation and conditions of [Theorem 3](#). Thus, first note that by [Theorem 4](#), A and B are regular since they are products of intervals. Also, $N_A(b, p, \rho) = \{0\}$ and $N_B(b, p, \rho)$ is the standard multiplier cone consisting of

Lagrange multipliers $(\lambda, \zeta, \tau, \phi) \in \mathbb{R}^{|D \times \Theta|} \times \mathbb{R} \times \mathbb{R}^{|\Theta|} \times \mathbb{R}^{|D \times \Theta|}$ with each multiplier being zero if an interior point, non-positive if a right endpoint only, non-negative if a left end-point only, and unrestricted if both right and left end-points of the corresponding interval of B .

We now show that the constraint qualification of Theorem 3 holds at (b, p, ρ) . To do this consider the matrix obtained from the Jacobian of F after eliminating the row corresponding to the gradient of the participation constraint function ($i = 2$). This matrix has the form

$$\begin{bmatrix} I & 0 \\ \tilde{C} & \tilde{D} \end{bmatrix}$$

where \tilde{D} is a square matrix and I is an identity matrix. Therefore, its determinant, therefore, is given by the product of the determinant of the identity matrix and the determinant of \tilde{D} . The matrix \tilde{D} has the form

$$\begin{bmatrix} \tilde{B} & 0 \\ \tilde{A} & -\tilde{B}' \end{bmatrix}.$$

where $\tilde{A} = \nabla^2 c(p)$ and $\tilde{B} = [I_{|\Theta| \times |\Theta|} \dots I_{|\Theta| \times |\Theta|}]$. It is enough to show that the matrix obtained from \tilde{D} by interchanging the first block and the second block of rows is invertible. Block-wise inversion of this matrix gives

$$\begin{bmatrix} (\tilde{B}\tilde{A}^{-1}\tilde{B}')^{-1} & -(\tilde{B}\tilde{A}^{-1}\tilde{B}')^{-1}\tilde{B}\tilde{A}^{-1} \\ -\tilde{A}^{-1}(-\tilde{B}')(\tilde{B}\tilde{A}^{-1}\tilde{B}')^{-1} & \tilde{A}^{-1} + \tilde{A}^{-1}(-\tilde{B}')(\tilde{B}\tilde{A}^{-1}\tilde{B}')^{-1}\tilde{B}\tilde{A}^{-1} \end{bmatrix}.$$

Hence it is invertible if \tilde{A} and $\tilde{B}\tilde{A}^{-1}\tilde{B}'$ are invertible. \tilde{A} is invertible because it is positive definite as the Hessian of a strictly convex function. To see that $\tilde{B}\tilde{A}^{-1}\tilde{B}'$ is invertible, let x be any nonzero vector conformable for pre-multiplication with \tilde{B} .

Then, since \tilde{A}^{-1} is positive definite,

$$0 < x' \tilde{B} \tilde{A}^{-1} \tilde{B}' x,$$

which implies that $\tilde{B} \tilde{A}^{-1} \tilde{B}'$ is positive definite and thus invertible. We can conclude from this that the set of gradients of the functions f_j^i ($i \neq 2$) at (b, p, ρ) form a linearly independent set of vectors.

At least one of the liability limits must not bind, and so we can replace the corresponding gradient with the gradient of f_j^2 . Since p is assumed to place positive probability on each decision in each state and the prior has full support, the first $|D \times \Theta|$ components of this gradient are nonzero. Thus the new set of vectors form a linearly independent set. Therefore the standard constraint qualification of Theorem 3 holds at (b, p, ρ) and hence C is regular at (b, p, ρ) (so that $\hat{N}_C(b, p, \rho) = N_C(b, p, \rho)$) and hence

$$N_C(b, p, \rho) = \left\{ \sum_{d, \theta} \lambda[d, \theta] \nabla f_{(d, \theta)}^1(b, p, \rho) + \zeta \nabla f^2(b, p, \rho) + \sum_{\theta} \tau[\theta] \nabla f_{\theta}^3(b, p, \rho) + \sum_{d, \theta} \phi[d, \theta] \nabla f_{(d, \theta)}^4(b, p, \rho) : (\lambda, \zeta, \tau, \phi) \in N_B(F(b, p, \rho)) \right\}.$$

The condition $\nabla_{(b, p, \rho)} E_p[y - b] \in \hat{N}_C(b, p, \rho)$ may now be written

$$\begin{aligned} \pi(\theta) p(d|\theta) &= \lambda[d, \theta] + \zeta \pi(\theta) p(d|\theta) + \phi[d, \theta] \pi(\theta), \\ \pi(\theta) (y(d, \theta) - b(d, \theta)) &= \tau[\theta] \\ + \sum_{d', \theta'} \phi[d', \theta'] \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')} + \zeta \left(\pi(\theta) b(d, \theta) - \frac{\partial c(p)}{\partial p(d|\theta)} \right), \\ 0 &= - \sum_{d'} \phi[d', \theta]. \end{aligned}$$

Therefore

$$b(d, \theta) = y(d, \theta) - \frac{\tau[\theta] + \xi\rho[\theta]}{\pi(\theta)} - \frac{1}{\pi(\theta)} \sum_{d', \theta'} \phi[d', \theta'] \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')}. \quad (\text{A.5})$$

where $\phi[d, \theta] = p(d|\theta)(1 - \xi) - \lambda[d, \theta]/\pi(\theta)$. Defining $\beta : \Theta \rightarrow \mathbb{R}$ and $\gamma : D \times \Theta \rightarrow \mathbb{R}$ as

$$\beta(\theta) = \frac{\tau[\theta] + \xi\rho[\theta]}{\pi(\theta)},$$

$$\gamma(d, \theta) = \frac{1}{\pi(\theta)} \sum_{d', \theta'} \left(p(d'|\theta')(1 - \xi) - \frac{\lambda[d', \theta']}{\pi(\theta')} \right) \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')}.$$

completes the proof. \square

A.1.4 Results used in the proofs of Propositions 1 and 3

Definition 4. (*Tangent vector*). A vector $w \in \mathbb{R}^n$ is tangent to a set $C \subseteq \mathbb{R}^n$ at a point $\bar{x} \in C$, written $w \in T_C(\bar{x})$, if there exists a sequence x_1, x_2, x_3, \dots in C converging to \bar{x} along with a sequence of positive scalars $\tau_1, \tau_2, \tau_3, \dots$ converging to 0, such that the sequence $(x_1 - \bar{x})/\tau_1, (x_2 - \bar{x})/\tau_2, (x_3 - \bar{x})/\tau_3, \dots$ converges to w .

Definition 5. (*Normal vector*). Let $C \subseteq \mathbb{R}^n$ and $\bar{x} \in C$. A vector v is normal to C at \bar{x} in the regular sense, written $v \in \hat{N}_C(\bar{x})$, if

$$v \cdot w \leq 0 \text{ for all } w \in T_C(\bar{x}).$$

It is normal to C (in the general sense), written $v \in N_C(\bar{x})$, if there are sequences x_1, x_2, x_3, \dots in C converging to \bar{x} , and v_1, v_2, v_3, \dots in $\hat{N}_C(x_1), \hat{N}_C(x_2), \hat{N}_C(x_3), \dots$ converging to v .

Definition 6. (*Clarke regularity of sets*). A closed set $C \subseteq \mathbb{R}^n$ is regular at one of its points \bar{x} in the sense of Clarke if every normal vector to C at \bar{x} is a regular normal vector, i.e. $N_C(\bar{x}) = \hat{N}_C(\bar{x})$.

Theorem 2. (6.12. of Rockafellar and Wets [47]) (Basic first order conditions for optimality). Consider a problem of maximizing a differentiable function f_0 over a set $C \subseteq \mathbb{R}^n$. A necessary condition for \bar{x} to be locally optimal is

$$\nabla f_0(\bar{x}) \in \hat{N}_C(\bar{x}).$$

When C is convex and f_0 concave, this condition is sufficient for \bar{x} to be globally optimal.

Theorem 3. (6.14 of Rockafellar and Wets [47]) (Normal cones to sets with a constraint structure). Let

$$C = \{x \in A : F(x) \in B\}$$

for closed sets $A \subseteq \mathbb{R}^n$ and $B \subseteq \mathbb{R}^m$ and a C^1 mapping $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$, written component-wise as $F(x) = (f_1(x), \dots, f_m(x))$. Suppose that the following assumption, to be called the standard constraint qualification at \bar{x} , is satisfied:

$$\left\{ \begin{array}{l} \text{the only vector } y \in N_B(F(\bar{x})) \text{ for which} \\ -(y_1 \nabla f_1(\bar{x}) + \dots + y_m \nabla f_m(\bar{x})) \in N_A(\bar{x}) \text{ is } y = (0, \dots, 0). \end{array} \right.$$

Then, if A is regular at \bar{x} and B is regular at $F(\bar{x})$, C is regular at \bar{x} and

$$N_C(\bar{x}) = \{y_1 \nabla f_1(\bar{x}) + \dots + y_m \nabla f_m(\bar{x}) + z : y \in N_B(F(\bar{x})), z \in N_A(\bar{x})\}.$$

Theorem 4. (6.15 of Rockafellar and Wets [47]) (Normals to boxes). Suppose $B = B_1 \times \dots \times B_n$, where each B_j is a closed interval in \mathbb{R} . Then B is regular at every one of its points $\bar{x} = (\bar{x}_1, \dots, \bar{x}_n)$, and its normal cones have the form $N_B(\bar{x}) = N_{B_1}(\bar{x}_1) \times \dots \times N_{B_n}(\bar{x}_n)$,

where

$$N_{B_j}(\bar{x}_j) = \begin{cases} [0, \infty) & \text{if } \bar{x}_j \text{ is only the left endpoint of } B_j, \\ (-\infty, 0] & \text{if } \bar{x}_j \text{ is only the right endpoint of } B_j, \\ \{0\} & \text{if } \bar{x}_j \text{ is an interior point of } B_j, \\ (-\infty, \infty) & \text{if } B_j \text{ is a one point interval.} \end{cases}$$

A.1.5 Derivation of the optimal contract in the multi-period model

The agent's first order condition now takes account of his continuation payoffs

$$\pi(\theta^t)(b(d^t, \theta^t) + \delta_A v_A(d^t, \theta^t)) - \frac{\partial c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta)} = \rho[d^{t-1}, \theta^t].$$

The principal's first order conditions are now

$$\begin{aligned} & \pi(\theta|\theta^{t-1})p(d|d^{t-1}, \theta^{t-1}, \theta) \\ &= \lambda[d^t, \theta^t] + \zeta[d^{t-1}, \theta^{t-1}]\pi(\theta|\theta^{t-1})p(d|d^{t-1}, \theta^{t-1}, \theta) + \phi[d^t, \theta^t]\pi(\theta|\theta^{t-1}), \\ & \quad \pi(\theta|\theta^{t-1})(y(d^t, \theta^t) + \delta_P v_P(d^t, \theta^t) - b(d^t, \theta^t)) = \tau[\theta^t, d^{t-1}] \\ & \quad + \sum_{d', \theta'} \phi[(d^{t-1}, d'), (\theta^{t-1}, \theta')] \frac{\partial^2 c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta) \partial p(d'|d^{t-1}, \theta^{t-1}, \theta')} \\ & \quad + \zeta[d^{t-1}, \theta^{t-1}] \left(\pi(\theta^t)(b(d^t, \theta^t) + \delta_A v_A(d^t, \theta^t)) - \frac{\partial c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta)} \right), \\ & \quad 0 = - \sum_{d'} \phi[(d^{t-1}, d'), \theta^t]. \end{aligned}$$

Combining these gives

$$\begin{aligned} & b(d^t, \theta^t) = y(d^t, \theta^t) + \delta_P v_P(d^t, \theta^t) - \frac{\tau[d^{t-1}, \theta^t] + \zeta \rho[d^{t-1}, \theta^t]}{\pi(\theta|\theta^{t-1})} \\ & - \frac{1}{\pi(\theta|\theta^{t-1})} \sum_{d', \theta'} \phi[(d^{t-1}, d'), (\theta^{t-1}, \theta')] \frac{\partial^2 c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta) \partial p(d'|d^{t-1}, \theta^{t-1}, \theta')}. \end{aligned}$$

where

$$\phi[(d^{t-1}, d'), (\theta^{t-1}, \theta')] = p(d|d^{t-1}, \theta^{t-1}, \theta)(1 - \zeta[d^{t-1}, \theta^{t-1}]) - \frac{\lambda[d^t, \theta^t]}{\pi(\theta|\theta^{t-1})}.$$

Defining

$$\beta(d^{t-1}, \theta^t) = \frac{\tau[d^{t-1}, \theta^t] + \zeta[d^{t-1}, \theta^{t-1}]\rho[d^{t-1}\theta^t]}{\pi(\theta|\theta^{t-1})},$$

$$\gamma(d^t, \theta^t) = \frac{1}{\pi(\theta|\theta^{t-1})} \sum_{d', \theta'} \phi[(d^{t-1}, d'), (\theta^{t-1}, \theta')] \frac{\partial^2 c(p|d^{t-1}, \theta^{t-1})}{\partial p(d|d^{t-1}, \theta^{t-1}, \theta) \partial p(d'|d^{t-1}, \theta^{t-1}, \theta')},$$

completes the derivation.

A.1.6 Posterior-separable costs and complementarities

Following their definition in [31] we suppose that the agent's cost function takes the form

$$c(p) = \sum_{d \in D} p(d) D(p(\cdot|d) || \pi)$$

where

$$D(p(\cdot|d) || \pi) = H(p(\cdot|d)) - H(\pi) - (p(\cdot|d) - \pi)^T \nabla H(\pi)$$

is the Bregman divergence associated with the convex function H and where

$$\frac{\partial}{\partial p(\check{\theta}|d)} \frac{\partial}{\partial p(\bar{\theta}|d)} H(p(\cdot|d)) = \frac{k(\check{\theta}, \bar{\theta})}{p(\check{\theta}|d)p(\bar{\theta}|d)}. \quad (\text{A.6})$$

We are interested in calculating

$$\frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d|\theta')}.$$

First note that

$$\frac{\partial}{\partial p(d|\theta')} p(\bar{\theta}|d) = \begin{cases} \frac{\pi(\theta')}{p(d)} (1 - p(\theta'|d)) & \text{if } \bar{\theta} = \theta' \\ -\frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) & \text{if } \bar{\theta} \neq \theta' \end{cases}.$$

We have

$$\begin{aligned}\frac{\partial c(p)}{\partial p(d|\theta')} &= \frac{\partial}{\partial p(d|\theta')} (p(d)D(p(\cdot|d)||\pi)) \\ &= \pi(\theta')D(p(\cdot|d)||\pi) + p(d)\frac{\partial}{\partial p(d|\theta')}D(p(\cdot|d)||\pi),\end{aligned}$$

and so

$$\begin{aligned}\frac{\partial^2 c(p)}{\partial p(d|\theta)\partial p(d|\theta')} &= \pi(\theta')\frac{\partial}{\partial p(d|\theta)}(D(p(\cdot|d)||\pi)) + \pi(\theta)\frac{\partial}{\partial p(d|\theta')}D(p(\cdot|d)||\pi) \\ &\quad + p(d)\frac{\partial}{\partial p(d|\theta)}\frac{\partial}{\partial p(d|\theta')} (D(p(\cdot|d)||\pi)).\end{aligned}\tag{A.7}$$

Calculating the parts of (A.7):

$$\begin{aligned}\frac{\partial}{\partial p(d|\theta')}D(p(\cdot|d)||\pi) &= \frac{\partial}{\partial p(d|\theta')} \left(H(p(\cdot|d)) - H(\pi) - (p(\cdot|d) - \pi)^T \nabla H(\pi) \right) \\ &= \frac{\partial}{\partial p(d|\theta')} H(p(\cdot|d)) - \frac{\partial}{\partial p(d|\theta')} \sum_{\bar{\theta}} (p(\bar{\theta}|d) - \pi(\bar{\theta})) \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})},\end{aligned}\tag{A.8}$$

where

$$\begin{aligned}\frac{\partial}{\partial p(d|\theta')} H(p(\cdot|d)) &= \sum_{\bar{\theta}} \frac{\partial}{\partial p(d|\theta')} (p(\bar{\theta}|d)) \frac{\partial}{\partial p(\bar{\theta}|d)} H(p(\cdot|d)) \\ &= \frac{\pi(\theta')}{p(d)} \frac{\partial}{\partial p(\theta'|d)} H(p(\cdot|d)) - \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \frac{\partial}{\partial p(\bar{\theta}|d)} H(p(\cdot|d)),\end{aligned}$$

and

$$\begin{aligned}\frac{\partial}{\partial p(d|\theta')} \sum_{\bar{\theta}} (p(\bar{\theta}|d) - \pi(\bar{\theta})) \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} &= \sum_{\bar{\theta}} \frac{\partial}{\partial p(d|\theta')} (p(\bar{\theta}|d)) \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \\ &= \frac{\pi(\theta')}{p(d)} \frac{\partial H(\pi)}{\partial \pi(\theta')} - \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})}.\end{aligned}\tag{A.9}$$

Putting these in (A.8)

$$\begin{aligned}\frac{\partial D(p(\cdot|d)||\pi)}{\partial p(d|\theta')} &= \frac{\pi(\theta')}{p(d)} \frac{\partial H(p(\cdot|d))}{\partial p(\theta'|d)} - \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} \\ &\quad - \frac{\pi(\theta')}{p(d)} \frac{\partial H(\pi)}{\partial \pi(\theta')} + \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})}\end{aligned}$$

$$\begin{aligned}
&= \frac{\pi(\theta')}{p(d)} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\theta'|d)} - \frac{\partial H(\pi)}{\partial \pi(\theta')} \right) - \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right) \\
&= \sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right).
\end{aligned}$$

Therefore

$$\begin{aligned}
\frac{\partial}{\partial p(d|\theta)} \frac{\partial D(p(\cdot|d)||\pi)}{\partial p(d|\theta')} &= \frac{\partial}{\partial p(d|\theta)} \sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right) \quad (\text{A.10}) \\
&= \sum_{\bar{\theta}} \frac{\partial^2 p(\bar{\theta}|d)}{\partial p(d|\theta) \partial p(d|\theta')} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right) + \sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \frac{\partial}{\partial p(d|\theta)} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} \right).
\end{aligned}$$

We have

$$\frac{\partial}{\partial p(d|\theta)} \frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} = \sum_{\check{\theta}} \frac{\partial p(\check{\theta}|d)}{\partial p(d|\theta)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\bar{\theta}|d)}.$$

Putting this into (A.10)

$$\begin{aligned}
&\frac{\partial}{\partial p(d|\theta)} \frac{\partial}{\partial p(d|\theta')} D(p(\cdot|d)||\pi) \\
&= \sum_{\bar{\theta}} \frac{\partial^2 p(\bar{\theta}|d)}{\partial p(d|\theta) \partial p(d|\theta')} \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right) + \sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \sum_{\check{\theta}} \frac{\partial p(\check{\theta}|d)}{\partial p(d|\theta)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\bar{\theta}|d)}.
\end{aligned}$$

Substituting into (A.7)

$$\begin{aligned}
&\frac{\partial c(p)}{\partial p(d|\theta) \partial p(d|\theta')} = \\
&\sum_{\bar{\theta}} \left(\pi(\theta') \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta)} + \pi(\theta) \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} + p(d) \frac{\partial^2 p(\bar{\theta}|d)}{\partial p(d|\theta) \partial p(d|\theta')} \right) \left(\frac{\partial H(p(\cdot|d))}{\partial p(\bar{\theta}|d)} - \frac{\partial H(\pi)}{\partial \pi(\bar{\theta})} \right) \\
&\quad + p(d) \sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \sum_{\check{\theta}} \frac{\partial p(\check{\theta}|d)}{\partial p(d|\theta)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\bar{\theta}|d)}.
\end{aligned}$$

We have

$$\begin{aligned}
&\sum_{\bar{\theta}} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} \sum_{\check{\theta}} \frac{\partial p(\check{\theta}|d)}{\partial p(d|\theta)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\bar{\theta}|d)} \\
&= \frac{\pi(\theta')}{p(d)} \left(\frac{\pi(\theta)}{p(d)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\theta|d) \partial p(\theta'|d)} - \sum_{\check{\theta}} \frac{\pi(\theta)}{p(d)} p(\check{\theta}|d) \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\theta'|d)} \right)
\end{aligned}$$

$$\begin{aligned}
& - \sum_{\bar{\theta}} \frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) \left(\frac{\pi(\theta)}{p(d)} \frac{\partial^2 H(p(\cdot|d))}{\partial p(\theta|d) \partial p(\bar{\theta}|d)} - \sum_{\check{\theta}} \frac{\pi(\theta)}{p(d)} p(\check{\theta}|d) \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\bar{\theta}|d)} \right) \\
& = \frac{\pi(\theta')\pi(\theta)}{p(d)^2} \left(R(d, \theta, \theta') - \sum_{\bar{\theta}} p(\bar{\theta}|d) R(d, \theta, \bar{\theta}) \right),
\end{aligned}$$

where

$$R(d, \theta, \theta') = \frac{\partial^2 H(p(\cdot|d))}{\partial p(\theta|d) \partial p(\theta'|d)} - \sum_{\check{\theta}} p(\check{\theta}|d) \frac{\partial^2 H(p(\cdot|d))}{\partial p(\check{\theta}|d) \partial p(\theta'|d)}. \quad (\text{A.11})$$

Note

$$\begin{aligned}
& \frac{\partial^2 p(\bar{\theta}|d)}{\partial p(d|\theta) \partial p(d|\theta')} = \frac{\partial}{\partial p(d|\theta)} \begin{cases} \frac{\pi(\theta')}{p(d)} (1 - p(\theta'|d)) & \text{if } \bar{\theta} = \theta' \\ -\frac{\pi(\theta')}{p(d)} p(\bar{\theta}|d) & \text{if } \bar{\theta} \neq \theta' \end{cases} \\
& = \begin{cases} -\frac{\pi(\theta')\pi(\theta)}{p(d)^2} (1 - p(\theta'|d)) - \frac{\pi(\theta')}{p(d)} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta)} & \text{if } \bar{\theta} = \theta' \\ -\frac{\pi(\theta')\pi(\theta)}{p(d)^2} p(\bar{\theta}|d) - \frac{\pi(\theta')}{p(d)} \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta)} & \text{if } \bar{\theta} \neq \theta' \end{cases} \\
& = \frac{\pi(\theta')\pi(\theta)}{p(d)^2} \begin{cases} -(1 - p(\theta'|d)) - (1 - p(\theta|d)) & \text{if } \bar{\theta} = \theta' \text{ and } \theta' = \theta \\ -(1 - p(\theta'|d)) + p(\theta'|d) & \text{if } \bar{\theta} = \theta' \text{ and } \theta' \neq \theta \\ p(\bar{\theta}|d) - (1 - p(\theta|d)) & \text{if } \bar{\theta} \neq \theta' \text{ and } \theta' = \theta \\ p(\bar{\theta}|d) + p(\theta'|d) & \text{if } \bar{\theta} \neq \theta' \text{ and } \theta' \neq \theta \end{cases}.
\end{aligned}$$

Therefore

$$\left(\pi(\theta') \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta)} + \pi(\theta) \frac{\partial p(\bar{\theta}|d)}{\partial p(d|\theta')} + p(d) \frac{\partial^2 p(\bar{\theta}|d)}{\partial p(d|\theta) \partial p(d|\theta')} \right) = 0.$$

Then

$$\frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d|\theta')} = \frac{\pi(\theta')\pi(\theta)}{p(d)} \left(R(d, \theta, \theta') - \sum_{\bar{\theta}} p(\bar{\theta}|d) R(d, \theta, \bar{\theta}) \right).$$

From (A.6) and (A.11)

$$R(d, \theta, \theta') = \frac{k(\theta, \theta')}{p(\theta|d)p(\theta'|d)} - \sum_{\check{\theta}} p(\check{\theta}|d) \frac{k(\check{\theta}, \theta')}{p(\check{\theta}|d)p(\theta'|d)} = \frac{k(\theta, \theta')}{p(\theta|d)p(\theta'|d)}$$

where the summing to zero of the columns and rows of the information cost matrix k has been used. Finally,

$$\frac{\partial^2 c(p)}{\partial p(d|\theta)\partial p(d|\theta')} = \frac{\pi(\theta')\pi(\theta)}{p(d)} \left(\frac{k(\theta, \theta')}{p(\theta|d)p(\theta'|d)} \right) = p(d) \frac{k(\theta, \theta')}{p(d|\theta)p(d|\theta')}. \quad (\text{A.12})$$

To check this expression, note that the inverse Fisher information matrix is

$$k(\theta, \theta') = \begin{cases} p(\theta|d)(1 - p(\theta|d)) & \text{if } \theta = \theta' \\ -p(\theta|d)p(\theta'|d) & \text{if } \theta \neq \theta' \end{cases}.$$

Substituting this into (A.12)

$$\frac{\partial c(p)}{\partial p(d|\theta)\partial p(d|\theta')} = \begin{cases} \frac{\pi(\theta')\pi(\theta)}{p(d)} \left(\frac{1}{p(\theta'|d)} - 1 \right) & \text{if } \theta = \theta' \\ -\frac{\pi(\theta')\pi(\theta)}{p(d)} & \text{if } \theta \neq \theta' \end{cases},$$

which is the Hessian of expected reduction in Shannon entropy. Proposition 3 gives the form of $\gamma : D \times \Theta \rightarrow \mathbb{R}$

$$\gamma(d, \theta) = \frac{1}{\pi(\theta)} \sum_{d', \theta'} \left(p(d'|\theta')(1 - \xi) - \frac{\lambda[d', \theta']}{\pi(\theta')} \right) \frac{\partial^2 c(p)}{\partial p(d|\theta)\partial p(d'|\theta')}.$$

Using (A.12) in this expression gives

$$\begin{aligned} \gamma(d, \theta) &= \frac{1}{\pi(\theta)} \sum_{\theta'} \left((1 - \xi) - \frac{\lambda[d, \theta']}{p(d|\theta')\pi(\theta')} \right) p(d) \frac{k(\theta, \theta')}{p(d|\theta)} \\ &= -\frac{1}{p(d)} \sum_{\theta'} \left(\frac{k(\theta, \theta')}{p(\theta|d)p(\theta'|d)} \right) \lambda[d, \theta']. \end{aligned}$$

A.1.7 Risk-aversion

Here we consider the case where the agent is risk-averse with a concave Bernoulli utility function u over wealth. When the agent is not capacity constrained the formula for the optimal contract extends straightforwardly. The only change is that

γ is now given by

$$\gamma(d, \theta) = \frac{1}{\pi(\theta)} \sum_{d', \theta'} \frac{1}{u'(b(d, \theta'))} \left(p(d'|\theta')(1 - u'(b(d, \theta')) \xi) - \frac{\lambda[d', \theta']}{\pi(\theta')} \right) \frac{\partial^2 c(p)}{\partial p(d|\theta) \partial p(d'|\theta')}.$$

For the cost functions of [31] this becomes

$$\gamma(d, \theta) = \sum_{\theta'} \left(\frac{k(\theta, \theta', p(\cdot|d))}{p(\theta|d)p(\theta'|d)} \right) \left(\frac{p(d, \theta') - \lambda[d, \theta']}{u'(b(d, \theta'))} \right).$$

In this case, γ takes on an insurance role as well as countering undesirable incentives for risk-taking induced by limited liability. It is not yet clear whether there is an analogous result to Proposition 2 when the agent is risk-averse.

A.2 Graphical method and an example

A.2.1 Graphical method

Here we use geometric methods to study the optimal contract when the agent's cost function is posterior separable. The agent's cost function is posterior-separable if there exists a nonnegative real valued function Y on $\Delta(\Theta)$ such that

$$c(p) = Y(\pi) - \sum_{d \in D} p(d)Y(p(\cdot|d)).$$

A benefit of using posterior separable cost functions is that they allow for a simple geometric representation of the agent's problem.¹ Figure A.1 represents the agent's problem of choosing an experiment under a posterior separable cost function when there are two states and two decisions.

Contract payments in states θ_1 and θ_2 are shown on the left and right vertical axes. The horizontal axis measures the probability assigned to state θ_2 . The dashed

¹This representation is taken from [10]; the Bayesian persuasion literature initiated by [37] develops a geometric approach similar to the one here using insights from [6].




Figure A.1: The agent's problem for $\pi(\theta_2) = 0.45$ and contract $b(d_1, \theta_1) = 0$, $b(d_2, \theta_1) = b(d_2, \theta_2) = 1$, and $b(d_1, \theta_2) = 2$. Reduced form $B(p)$ and upper envelope of net utilities $B(p) + Y(p)$ as a function of the posterior probability assigned to state θ_2 . Optimal posteriors are given by maximizing expected net utility subject to the average posterior equaling the prior.

lines connecting contract payments in different states show the expected payoffs of each decision as a function of the probability assigned to state θ_2 . The upper envelope of these lines is called the *reduced form* of contract b and is denoted by $B(p)$ (That is, $B : \Delta(\Theta) \rightarrow \mathbb{R}$ is defined as the agent's maximal expected contract payment as a function of his posterior belief about the state: $B(p) = \max\{E_p[b(d, \cdot)] : d \in D\}$). The reduced form describes the decision the agent will take as a function of his posterior belief about the state. For the contract shown, the agent will choose decision d_2 when the probability he assigns to state θ_2 is less than $1/2$ and he will choose decision d_1 when the probability he assigns to state θ_2 is greater than $1/2$.

Because the cost function is posterior separable the agent's objective function may be written as

$$\sum_d p(d) \left(\sum_\theta p(\theta|d) b(d, \theta) + Y(p(\cdot|d)) \right) - Y(\pi).$$

Following [10], the term $\sum_{\theta} p(\theta|d)b(d, \theta) + Y(p(\cdot|d))$ is called the *net utility* of decision d . The upper envelope of the net utility for decision d_1 and the net utility for decision d_2 is given by $B(p) + Y(p)$. Note that the average of the agent’s posterior beliefs is equal to his prior belief and therefore the agent’s utility is given by weighting the optimal net utilities—given by $B + Y$ —with weights that average the posteriors to the prior. In other words, the optimal posteriors are computed by considering the “best” tangent hyperplane to the upper envelope of the net utility function, $B + Y$, or, alternatively, as in [37], by evaluating the concavification of $B + Y$ at the prior.² In the figure, the optimal posterior probabilities are shown as well as the point $(\pi(\theta_2), v_A(b) + Y(\pi))$ corresponding (but not equal) to the agent’s maximized utility.

Altering a contract by a state-dependent transfer. The agent’s behavior is unchanged when his contract is altered by a state-dependent transfer since $E_p[b - \beta] - c(p) = E_p[b] - E_{\pi}[\beta] - c(p)$ for all $\beta : \Theta \rightarrow \mathbb{R}$. Figure A.2 illustrates this for the contract shown in black and the state contingent payment $\beta(\theta_1) = 0, \beta(\theta_2) = 1$. Note that the optimal experiment would not change if $\beta(\theta_2) > 1$ but the resulting contract would not be feasible.

²The *concavification* of a function is the smallest concave function everywhere weakly greater than the function.






Figure A.2: Altering the contract b by the state contingent payment $\beta(\theta_1) = 0, \beta(\theta_2) = 1$ does not alter the experiment the agent chooses.

Scaling a contract when capacity constraints binds. When the agent's capacity constraint binds under his chosen experiment it is possible to scale down—multiply by a factor less than one—the contract without altering the agent's behavior. Figure A.3a shows that imposing a binding capacity constraint alters the agent's chosen experiment by bringing the posterior probabilities closer to the prior. Figure A.3b shows that given the binding capacity constraint the contract b can be scaled down by any factor α in $[1/2, 1]$ without altering the agent's chosen experiment.



(a)



(b)

Figure A.3: (a) A binding capacity constraint (thick dashed lines) brings the posterior probabilities closer to the prior. (b) With a binding capacity constraint (thick dashed lines), scaling incentives in half (thick solid lines) does not alter the agent's chosen experiment.

A.2.2 A two-state two-decision example

Consider the following problem in which a principal hires an agent to choose between a safe and risky course of action:

$y(d, \theta)$	θ_1	θ_2
d_1	0	10
d_2	5	5

Suppose the agent's prior belief on state θ_1 is $2/3$ and that the set of feasible experiments are those with cost less than $k = 1/2$.

First-best contracts

We first solve for Pareto optimal contracts that induce first-best experimentation. Since altering a contract by a state-dependent transfer does not alter the experiment chosen by the agent, any contract $y - \beta$ with $\beta : \Theta \rightarrow \mathbb{R}$ induces first best experimentation. Similarly, if the agent's capacity constraint binds under these contracts, then there exists a number α' in $[0, 1]$ (as defined in Proposition 1) such that for each of the contracts $\alpha(y - \beta)$ with α in $[\alpha', 1]$ the agent chooses the same first-best experiment. Therefore, each contract of the form

$$\alpha(y - \beta) \text{ with } \alpha \in [\alpha', 1] \text{ and } \beta : \Theta \rightarrow \mathbb{R} \text{ chosen so the contract is feasible} \quad (\text{A.13})$$

is Pareto optimal and induces first-best experimentation. The least favorable of these contracts for the agent is $\alpha'(y + \beta')$ with β' defined by $\beta'(\theta_1) = 0$, $\beta'(\theta_2) = -5$. So the set of agent reservation utilities corresponding to these contracts is $[v_A(\alpha'(y - \beta')), v_A(y)]$.

Results from [41] and [10] show that when the agent's cost function is expected Shannon entropy reduction the optimal experiment resembles a logit-rule with the observable part of the utility from taking decision d in state θ given by $b(d, \theta)/(1 +$

$\mu) + \log p(d)$:

$$p(d|\theta) = \frac{p(d)e^{\frac{b(d,\theta)}{1+\mu}}}{\sum_{\bar{d}} p(\bar{d})e^{\frac{b(\bar{d},\theta)}{1+\mu}}}. \quad (\text{A.14})$$

The term μ is a nonnegative dual variable which is chosen as small as possible subject to the agent's capacity constraint binding. Using Equation (A.14), if the agent's capacity constraint were not to bind then the first-best experiment would be as shown in Table A.1a. However, this experiment is not feasible since its cost

$p(d \theta)$	θ_1	θ_2	$p(d \theta)$	θ_1	θ_2
d_1	0.007	0.993	d_1	0.031	0.969
d_2	0.993	0.007	d_2	0.969	0.031
(a)			(b)		

Table A.1: (a) First best experiment when the agent's capacity constraint does not bind. (b) First best experiment under capacity constraint $c(p) \leq 1/2$.

is 0.596. Instead, the agent chooses the experiment show in in Table A.1b which has cost 1/2 and is given by setting $\mu = 0.446$ in Equation (A.14) with $b = y$. This implies that $\alpha' = 1/(1 + \mu) = 0.692$. Under this experiment, the agent is slightly more prone to error than under the unconstrained experiment.

The minimum agent utility for a first-best contract is $v_A(\alpha'(y - \beta')) = 2.853$ and the maximum utility is $v_A(y) = 6.014$; the first-best Pareto optimal contracts given by Equation (A.13) deliver the agent utilities in this range. Under each of these contracts, the agent chooses the experiment shown in Table A.1b. Note that any Pareto optimal contract that delivers the agent a utility in this range must induce first best experimentation (otherwise the contracts found here would be an improvement, contradicting Pareto optimality).

Second-best contracts

In the previous section, we showed that if the agent's reservation utility is above 2.853, then all Pareto optimal contracts induce first-best experimentation. We now consider the case where the agent's reservation utility is below this level.

For a second-best contract the minimum payment in each state is 0. Consider the case where $b(d_1, \theta_1)$ and $b(d_2, \theta_2)$ are 0 and $b(d_1, \theta_2)$ and $b(d_2, \theta_1)$ are positive and such that the principal's liability limits do not bind. Then $\lambda[d_1, \theta_2]$ and $\lambda[d_2, \theta_1]$ are 0. Using the form of contract given in Proposition 3 and section 1.3.5 as well as the formula for the state-dependent transfer gives

$$\begin{aligned}\hat{\gamma}(d_1) &= \frac{\lambda[d_1, \theta_1]}{p(d_1)}, \beta(\theta_1) = \alpha y(d_1, \theta_1) + \hat{\gamma}(d_1) + \frac{\lambda[d_1, \theta_1]}{p(d_1, \theta_1)} \\ \hat{\gamma}(d_2) &= \frac{\lambda[d_2, \theta_2]}{p(d_2)}, \beta(\theta_2) = \alpha y(d_2, \theta_2) + \hat{\gamma}(d_2) + \frac{\lambda[d_2, \theta_2]}{p(d_2, \theta_2)}\end{aligned}$$

Since $\sum_d \lambda[d, \theta] = (1 - \xi)\pi(\theta)$, we have $\lambda[d_1, \theta_1] = (1 - \xi)\pi(\theta_1)$ and $\lambda[d_2, \theta_2] = (1 - \xi)\pi(\theta_2)$ and so

$$b(d_1, \theta_2) = \alpha(10 - 5) + \frac{(1 - \xi)\pi(\theta_2)}{p(d_2)} - \frac{(1 - \xi)\pi(\theta_2)}{p(d_2, \theta_2)} - \frac{(1 - \xi)\pi(\theta_1)}{p(d_1)} \quad (\text{A.15})$$

$$b(d_2, \theta_1) = \alpha(5 - 0) + \frac{(1 - \xi)\pi(\theta_1)}{p(d_1)} - \frac{(1 - \xi)\pi(\theta_1)}{p(d_1, \theta_1)} - \frac{(1 - \xi)\pi(\theta_2)}{p(d_2)} \quad (\text{A.16})$$

Recall that the minimum agent reservation utility such that the optimal contract induces first-best experimentation is 2.853. For this reservation utility the unique optimal contract is given by

$$\frac{1}{1 + \mu}(y + \beta'),$$

where $\mu = 0.446$ and $\beta'(\theta_1) = 0$, $\beta'(\theta_2) = 5$. Note that for a first-best contract we have $\xi = 1$ since a fall in the agent's reservation utility induces a one-for-one increase in the principal's utility. Using $\xi = 1$ in Equations (A.15) and (A.16) and setting α

such that the cost of the optimal experiment is $k = 0.5$, i.e. $\alpha = 1/(1 + 0.446)$, gives the same contract as the first-best contract for reservation utility 2.853. To solve for this optimal contract in general requires solving the system of equations (A.14), (A.15), and (A.16) and the inequalities given by the agent's capacity constraint and participation constraint.

A special case arises when the prior is uniform. Then a solution to the system of equations is given by $p(d_1) = p(d_2) = 1/2$, $p(d_1, \theta_1) = p(d_2, \theta_2)$, $b(d_1, \theta_2) = b(d_2, \theta_1)$ which implies that the optimal contract is affine and given by $b(d_1, \theta_2) = 5\alpha - t$, $b(d_2, \theta_1) = 5\alpha - t$, and 0 otherwise, where $t = (1 - \xi)/p(d_1|\theta_1) = (1 - \xi)/p(d_2|\theta_2)$.

We now develop some graphical intuition for the second-best Pareto optimal contracts characterized in this example making the simplifying assumptions that the agent's capacity constraint does not bind—so that $\alpha^* = 1$ —and that the agent's reservation utility is sufficiently low that his participation constraint does not bind—which implies $\xi = 0$. Solving Equations (A.14), (A.15), and (A.16) then gives the following optimal contract, experiment, state-dependent transfer β , and decision-dependent transfer γ

$b = y - \beta - \gamma$	θ_1	θ_2
d_1	0	1.00
d_2	0.702	0

(a)

$p(d \theta)$	θ_1	θ_2
d_1	0.160	0.514
d_2	0.840	0.486

(b)

β	θ_1	θ_2
d_1	3.836	6.596
d_2	3.836	6.596

(c)

γ	θ_1	θ_2
d_1	-3.836	2.404
d_2	0.462	-1.596

(d)

Table A.2: (a) Optimal contract (b) Optimal experiment (c) State-dependent transfer (d) Optimal distortion

We will now build this optimal contract in stages. First we consider the contract $y - \beta$ shown in Table A.3a. Since this contract is output adjusted by a state-dependent transfer the agent chooses a first-best experiment. Table A.3b presents the optimal experiment for this contract and Figure A.4 depicts the optimal posterior probabilities.

$y - \beta$	θ_1	θ_2
d_1	-3.836	3.404
d_2	1.164	-1.596

(a)

$p(d \theta)$	θ_1	θ_2
d_1	0.007	0.993
d_2	0.993	0.007

(b)

Table A.3: (a) Contract $y - \beta$. (b) First best experiment under contract $y - \beta$.




Figure A.4: Contract $y - \beta$ induces first best experimentation.

The contract $y - \beta$ induces very precise experimentation with the agent choosing the wrong decision with a 0.7 percent probability in either state.

Now consider the feasible contract induced by $y - \beta$ which is its truncation $\max\{0, y - \beta\}$. Table A.4 shows the optimal experiment for this contract and Figure A.5 depicts the optimal posterior probabilities (thick lines) and compares it to the optimal posterior probabilities for $y - \beta$ (thin lines)

$p(d \theta)$	θ_1	θ_2
d_1	0.211	0.963
d_2	0.789	0.037

Table A.4: Optimal experiment for contract $\max\{0, y - \beta\}$




Figure A.5: Optimal experiments for $\max\{0, y - \beta\}$ (thick lines) and $y - \beta$ (thin lines).

Under contract $\max\{0, y - \beta\}$, when the agent chooses decision d_2 he does so with approximately the same level of confidence in which state it is as under the first-best experiment. But when the agent chooses decision d_1 he does so with much less confidence about the state than under the first-best experiment.

The optimal experiment for contract $\max\{0, y - \beta\}$ shown in Table A.4 shows that in state θ_2 the agent chooses the incorrect experiment d_2 with a similar probability as under the first best experiment. However the agent now incorrectly chooses decision d_1 in state θ_1 about 20 percent of the time compared to 2 percent under the first best experiment. This effect arises because truncating the contract $y - \beta$ eliminates a greater downside for decision d_1 than d_2 and so encourages the agent to choose d_1 .

A way to measure the change in the principal's utility due to truncating the

contract $y - \beta$ is shown in Figure A.6. Point A gives the expected value of output less the constant $E_\pi[\beta(\theta)]$ and point B gives the expected contract payment. Since under the contract $y - \beta$ these points coincide and now point A is below B, the move from contract $y - \beta$ to $\max\{0, y - \beta\}$ reduces the principal's utility.




Figure A.6: Point A gives expected output minus $E_\pi[\beta(\theta)]$. Point B gives the expected contract payment.

Figure A.7 shows the optimal posterior beliefs for the optimal contract $b = y - \beta - \gamma$. Table A.2b gives the optimal experiment and Table A.5 lists the expected value of output, the contract, and the experiment costs under the three different contracts.

	$E_p[y]$	$E_p[b]$	$c(p)$
$y - \beta$	6.633	1.877	0.596
$\max\{0, y - \beta\}$	5.900	1.704	0.293
$b = y - \beta - \gamma$	5.321	0.566	0.067

Table A.5: Comparison of expected payoffs and costs




Figure A.7: The optimal contract. The lines defining A and B coincide so the optimal contract $b = y - \beta - \gamma$ has the same expected value as the contract $y - \beta$ at the optimal posteriors.

You can see that this contract is locally optimal since the lines defining A and B coincide and so for small perturbations the change in the expected contract payment will be roughly equal to the change in expected output. This also implies that the expected value of the optimal contract at the optimal posterior beliefs for this contract coincide with the expected value of the first best contract $y - \beta$ at these

posterior beliefs.

Figure A.8 shows a perturbation away from the optimal contract. You can see that the expected contract payment falls less than expected output so the perturbation is not profitable for the principal.




Figure A.8: A perturbation (thick line) away from the optimal contract. The expected contract payment (middle dot) falls less than expected output (lower dot).

Appendix B

Appendix to Chapter 2

B.1 Risk-aversion

We have been unable to find an analogue of our result for when the agent is risk-averse. If the principal is risk-neutral and the agent is risk-averse the problem in our proof is that [A.3](#) now involves a Bernoulli utility function u and therefore we cannot conclude (i) $E_p[u(b)] - E_{p_\alpha}[u(b_\alpha)] = 0$ from [A.2](#), nor (ii) that (i) implies $E_p[u(b)] - E_{p_\alpha}[u(b_\alpha)] = 0$.

The problem may be further seen by considering first-order conditions for the simple capacity constrained problem

$$\max_{b,p} E_p[y - b]$$

such that $c(p) \leq k$,

$$E_p[u(b)] - c(p) \geq E_{p'}[u(b)] \text{ for all } p' \text{ such that } c(p') \leq k, \quad (\text{B.1})$$

and

$$E_p[u(b)] - c(p) \geq 0.$$

Here the set of feasible contracts is unrestricted and the set of feasible distributions are those satisfying the capacity constraint $c(p) \leq k$. The agent's program may be written

$$\max_p E_p[u(b)] - c(p)$$

such that $p(\omega) \geq 0$ for all $\omega \in \Omega$, $\sum_{\omega \in \Omega} p(\omega) = 1$, and $c(p) \leq k$. If the agent's cost function is convex and the optimal distribution always assigns positive probability to each state, then a necessary and sufficient condition for p to be an optimal contract is

$$u(b(\omega)) - \frac{\partial c(p)}{\partial p(\omega)} = \rho + \mu \frac{\partial c(p)}{\partial p(\omega)}, \quad (\text{B.2})$$

where ρ and μ are Lagrange multipliers on the constraints $\sum_{\omega \in \Omega} p(\omega) = 1$ and $c(p) \leq k$. If the agent's cost function is strictly convex, then one may show that the program obtained after replacing the incentive constraint [B.1](#) by the first order-condition [B.2](#) and making μ a choice variable satisfies a qualification constraint which implies that any solution (in which p assigns positive probability to each state) satisfies first-order conditions

$$\begin{aligned} y(\omega) - b(\omega) &= \tau + \delta \frac{\partial c(p)}{\partial p(\omega)} - (\mu + 1) \sum_{\omega' \in \Omega} \phi[\omega'] \frac{\partial^2 c(p)}{\partial p(\omega) \partial p(\omega')} + \xi \left(u(b(\omega)) - \frac{\partial c(p)}{\partial p(\omega)} \right) \\ -p(\omega) &= \phi[\omega] u'(b(\omega)) + \xi p(\omega) u'(b(\omega)) \\ 0 &= - \sum_{\omega \in \Omega} \phi[\omega] \frac{\partial c(p)}{\partial p(\omega)} \end{aligned}$$

where τ and δ are Lagrange multipliers on the adding up and capacity constraint $\sum_{\omega \in \Omega} p(\omega) = 1$ and $c(p) \leq k$, $\phi[\omega]$ is a Lagrange multiplier on agent's first order condition [B.2](#), and ξ is a Lagrange multiplier on the participation constraint $E_p[b] - c(p) \geq 0$. Rearranging the first conditions and using the agent's first-order

condition B.2 gives

$$b(\omega) + \frac{\delta + \mu}{1 + \mu} u(b(\omega)) = y(\omega) - A - B \sum_{\omega' \in \Omega} p(\omega) \left(\frac{1}{u'(b(\omega))} + \xi \right) \frac{\partial^2 c(p)}{\partial p(\omega) \partial p(\omega')}$$

for some constants A and B . If the agent is risk-neutral, then this may be written as

$$b(\omega) = \frac{1 + \mu}{1 + \mu + \mu\delta} y(\omega) - A - B \sum_{\omega' \in \Omega} p(\omega) \left(\frac{1}{u'(b(\omega))} + \xi \right) \frac{\partial^2 c(p)}{\partial p(\omega) \partial p(\omega')},$$

which suggests our result. Such a representation does not seem to be possible when the agent is risk-averse.

B.2 Discounting

Discounting does not pose a problem when contract payoffs arise at a single date. In this case we can define a new cost function for the agent equal to his old cost function divided by his discount factor. However, when payoffs arrive at different dates our proof fails for similar reasons as when the agent is averse to risk. Suppose that payments can be made at two dates: date 0 and date 1. Suppose that the principal and agent discount payments with the principal's discount factor denoted by $\delta_P \in [0, 1]$ and the agent's discount factor denoted by $\delta_A \in [0, 1]$. Output and contracts are now functions of dates as well as states: $y, b : \{0, 1\} \times \Omega \rightarrow \mathbb{R}$. The principal's utility from profile (b, p) is now

$$E_{p(\omega)} \left[\sum_{t \in \{0, 1\}} \delta_P^t (y(t, \omega) - b(t, \omega)) \right]$$

and the agent's utility is

$$E_{p(\omega)} \left[\sum_{t \in \{0, 1\}} \delta_A^t b(t, \omega) \right] - c(p).$$

Then following the proof of the case without discounting we obtain the inequalities

$$\begin{aligned}
& E_{p(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t y(t, \omega) \right] - E_{p_\alpha(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t y(t, \omega) \right] \\
& \geq E_{p(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t b(t, \omega) \right] - E_{p_\alpha(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t b_\alpha(t, \omega) \right] \tag{B.3}
\end{aligned}$$

$$\geq \alpha \left(E_{p(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t y(t, \omega) \right] - E_{p_\alpha(\omega)} \left[\sum_{t \in \{0,1\}} \delta_P^t y(t, \omega) \right] \right) \geq 0,$$

$$c(p) - c(p_\alpha) \geq E_{p(\omega)} \left[\sum_{t \in \{0,1\}} \delta_A^t b(t, \omega) \right] - E_{p_\alpha(\omega)} \left[\sum_{t \in \{0,1\}} \delta_A^t b_\alpha(t, \omega) \right]. \tag{B.4}$$

As under risk-aversion, we cannot conclude that the value of discounted contract payments for the agent is the same under both contracts, nor that this implies their discounted value is equal for the principal.

Appendix C

Appendix to Chapter 3

C.1 Perfect competition

This appendix reviews the classic core equivalence papers of [19] and [5]. [19] prove that the core converges to the set of competitive allocations as the number of consumers increases by replication. [5] proves that core equivalence holds exactly if the set of consumers is modelled as an atomless measure space. Aumann's model is the canonical model of perfect competition so it is given a thorough treatment.

C.1.1 The Arrow-Debreu economy

In this section we define an Arrow-Debreu economy, the notions of core and competitive equilibrium, and the idea of perfect competition.

An exchange economy consists of commodities and consumers. Each commodity is defined by its physical characteristics and the location, date, and state of the world in which it is available. Each consumer is defined by his tastes over commodities and his initial holding of commodities. There is no production. All resources are privately owned.

An Arrow-Debreu exchange economy has a finite number of commodities and consumers. There are m consumers and ℓ commodities. Each consumer is indexed by i running from 1 to m and each commodity is indexed by h running from 1 to ℓ . \mathbb{R}^ℓ is called the *commodity space*. A *commodity bundle* is an element of the commodity space.

An *allocation* is an m -tuple $x = (x_1, \dots, x_m)$ of commodity bundles, where x_i represents the commodity bundle of consumer i . There is a fixed allocation, denoted $\omega = (\omega_1, \dots, \omega_m)$, called the *initial allocation*; ω_i is called the *initial endowment* of consumer i . An allocation x is said to be *feasible* if $\sum_{i=1}^m x_i = \sum_{i=1}^m \omega_i$.

A *preference relation*, denoted \succeq , is a complete preordering¹ over the commodity space and induces a *strict preference relation* (denoted \succ) in the following way: If x and y are arbitrary commodity bundles, then

$$x \succ y \text{ if and only if } x \succeq y \text{ and not } y \succeq x.$$

$x \succeq y$ is interpreted to mean “ x is preferred or indifferent to y .”

Definition 7. An Arrow-Debreu economy, denoted \mathcal{E} , is a triad $(\{1, \dots, m\}, \omega, \{\succeq_i\}_{i=1}^m)$. It consists of a set of consumers $\{1, \dots, m\}$, an initial allocation $\omega = (\omega_1, \dots, \omega_m)$, and a preference relation \succeq_i for each consumer $i = 1, \dots, m$.

Competitive equilibrium

For each commodity there is a price corresponding to the amount that must be paid now for one unit of that commodity. A consumer considers his wealth to be the value of his initial endowment at prevailing prices and demands the best commodity bundle that he can afford judged by his preferences. The market operates

¹a complete, transitive, and reflexive binary relation.

by adjusting prices to the point where aggregate demand meets aggregate supply. Trade then takes place, resulting in an allocation. An allocation generated in this way is called a *competitive allocation*.

The attractiveness of this process is its simplicity. Each agent only need know his own preferences, endowment, and the prevailing prices. He does not need to worry about the resources or preferences of his fellow consumers – prices summarise this information, representing aggregate preferences and scarcity.

Formally, a *price system*, denoted p , is a vector in \mathbb{R}_+^ℓ .

Definition 8. A competitive equilibrium (p, x) is a price system p and an allocation x such that x is feasible and for $i = 1, \dots, m$

1. $p \cdot x_i \leq p \cdot \omega_i$ and
2. $y \succ_i x_i$ implies $p \cdot y > p \cdot \omega_i$.

If (p, x) is a competitive equilibrium, then p is called an *equilibrium price system* and x a *competitive allocation*.

Perfect competition

The simplicity of our model of competition does not come for free. In order for the competitive allocation to make sense consumers must be *price takers*. This means that at each price consumers calculate their wealth according to prices and demand the best commodity bundle they can afford. The trouble is that consumers may be able to manipulate prices by, for example, throwing away some of their initial endowment.

Perfect competition solves this problem. It characterises a market in which price taking is endogenous. In a perfectly competitive market it is always optimal for consumers to be price takers. The idea behind perfect competition is to make the

market for each commodity so large that each consumer owns a *negligible* fraction of total resources. In turn, the influence of a single consumer becomes negligible.

With a finite number of commodities each consumer can be made negligible by increasing the number of consumers in the market for each commodity. It is important to note a finite number of commodities is required. If the economy has an infinite number of commodities, then even with an infinite number of consumers the market for each commodity could be small.

Of course, even in economies where consumers are negligible certain groups may collude together, pooling their resources to gain influence over prices. The idea we have in mind here is a cartel or a trade union. There is no concrete condition that will eliminate this possibility. We simply assume that consumers do not collude: that they are *independent* of one another.

We define a perfectly competitive exchange economy to be one in which each consumer is *negligible* and behaves *independently* of other consumers. By *negligible* we mean that each consumer owns a negligible fraction of total resources. By *independent* we mean that consumers do not collude with one another.

The core

We can describe an efficient allocation for an exchange economy by considering the cooperative outcome: the allocation arising from contract. This gives rise to the notion of the core – the set of allocations such that no non-negligible group of consumers can contract to redistribute their initial endowments so each is better off.

Formally, we define a *coalition* to be any nonempty set of consumers. An allocation x is *blocked* by a coalition S if there exists an allocation y such that $\sum_{i \in S} y_i = \sum_{i \in S} \omega_i$ and $y_i \succ_i x_i$ for each i in S .

Definition 9. *The core is the set of feasible allocations not blocked by any coalition.*

Given the simplicity of the competitive equilibrium we might question whether there is anything remarkable about it. There is no immediate reason to expect the competitive allocation to have special welfare properties. However it turns out that a competitive allocation is in the core. An allocation generated by competition is a solution to the cooperative problem.

The logic behind this result is the following: If a competitive allocation were not in the core, then some coalition could block it. Each member of the coalition would improve his situation by exchanging for a preferred bundle. By the definition of a competitive equilibrium no such commodity bundle is affordable at equilibrium prices. This implies the absurd conclusion that the total value of the coalition's commodity bundles is greater than the total value of their initial endowments.

Theorem 5. *A competitive allocation is in the core.*

Proof. Let x be a competitive allocation and p the equilibrium price system for x . Suppose by contradiction that x does not belong to the core. Then a coalition S blocks x with y . Because $y_i \succ_i x_i$ we have $p \cdot y_i > p \cdot \omega_i$ for each i in S . Hence

$$p \cdot \sum_{i \in S} y_i > p \cdot \sum_{i \in S} \omega_i,$$

contradicting $\sum_{i \in S} y_i = \sum_{i \in S} \omega_i$. □

C.1.2 Debreu and Scarf: A limit theorem on the core of an economy

[19] consider an Arrow-Debreu economy and a sequence of 'replications' of this economy. They show that as the number of replications tends to infinity the core converges to the set of competitive allocations.

The mathematical model

Debreu and Scarf consider the following mathematical model. Let \mathcal{E} be an Arrow-Debreu economy in which the following holds.

1. *Monotonicity*: let x and y be arbitrary commodity bundles. Then² $x > y$ implies $x \succ_i y$.
2. *Continuity*: let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_i z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_i x\}$ are open.
3. *Strong convexity*: let x and y be arbitrary distinct commodity bundles with $y \succeq_i x$ and let $0 < \alpha < 1$. Then $\alpha x + (1 - \alpha)y \succ_i y$.
4. ω_i is strictly positive in each of its coordinates.

The core as the number of consumers becomes infinite

Debreu and Scarf show that the core converges to the set of competitive allocations as the number of consumers becomes infinite. The main conceptual difficulty in proving this is that the core and competitive allocations are subsets of a space whose dimension depends on the number of consumers. Consequently it is difficult to define what is meant by convergence. Debreu and Scarf sidestep this problem by increasing the number of consumers by replication; starting with the economy \mathcal{E} they define the r -replica economy \mathcal{E}^r to be the economy consisting of rm consumers indexed by iq where $i = 1, \dots, m$ and $q = 1, \dots, r$. i refers to the type of consumer (i.e. those belonging to the original economy) and q refers to different consumers of the same type.

²Let x and y be arbitrary points in \mathbb{R}^ℓ . $x > y$ if and only if $x \neq y$ and $x_i \geq y_i$ for all $i = 1, \dots, \ell$.

The reason for considering replica economies is that core allocations assign the same commodity bundle to consumers of the same type. This is called the *equal treatment property*. Thus a core allocation

$$x = ((x_{11}, \dots, x_{1r}), (x_{21}, \dots, x_{2r}), \dots, (x_{m1}, \dots, x_{mr}))$$

for an r -replica economy can be summarised $x = (x_1, \dots, x_m)$. Since a competitive allocation is a core allocation this means that both core and competitive allocations can be analysed in $\mathbb{R}^{\ell m}$, even while varying r .

The intuition of the equal treatment property is this: Consider a core allocation and suppose it does not assign the same commodity bundle to consumers of the same type. Consider the coalition consisting of one consumer of each type who gets the least desired commodity bundle among the members of his type. Now assign each of these consumers the average commodity bundle of their type. This is feasible for the coalition because the core allocation is feasible and each member of each type has the same initial endowment. By strong convexity of preferences and our assumption, at least one member of the coalition strictly prefers this commodity bundle to the one he gets in the core allocation. By monotonicity and continuity of preferences this consumer can make each member of the coalition better off by giving a sufficiently small fraction of his commodity bundle away. Thus the coalition blocks the core allocation which is nonsense.

Formally,

Lemma 6. *An allocation in the core of \mathcal{E}^r assigns the same commodity bundle to all consumers of the same type.*

Proof. Let x be a core allocation for \mathcal{E}^r . Suppose by contradiction that x does not assign the same commodity bundle to consumers of the same type. Relabeling if necessary we may assume that consumer $i1$ receives the least desired commodity

bundle among the consumers of his type; that is $x_{iq} \succeq_i x_{i1}$ for all $i = 1, \dots, m$. Consider the coalition S consisting of consumers $\{1, \dots, m\}$. The allocation which gives $\frac{1}{r} \sum_{q=1}^r x_{iq}$ to each consumer i is feasible since

$$\sum_{i=1}^m \frac{1}{r} \sum_{q=1}^r x_{iq} = \frac{1}{r} \sum_{i=1}^m r \omega_i = \sum_{i=1}^m \omega_i.$$

For at least one consumer of type i in the coalition our assumption implies $x_{i1} \neq x_{iq}$ for some $q = 2, \dots, r$. Denote this type i^* . By strong convexity

$$\frac{1}{r} \sum_{q=1}^r x_{iq} \succeq_i x_{i1}$$

for all $i = 1, 2, \dots, m$, and for i^*

$$\frac{1}{r} \sum_{q=1}^r x_{i^*q} \succ_{i^*} x_{i^*1}.$$

By monotonicity and continuity, if the consumer of type i^* in the coalition gives a sufficiently small amount of his commodity bundle to the other members of the coalition, then each member of the coalition prefers the resulting commodity bundle to the commodity bundle obtained under x . Hence S blocks x . A contradiction. \square

Notice that the cores of replica economies form a decreasing sequence in r since a coalition of an r replica economy is a coalition of an $r + 1$ replica economy. Also notice that the set of competitive allocations is fixed as r varies. Because the number of coalitions increases as we increase the number of replications there is some hope that the core will converge to the set of competitive allocations as the number of replications tends to infinity. This is indeed the case. Debreu and Scarf prove

Theorem 6. *If $x = (x_1, \dots, x_m)$ is in the core of \mathcal{E}^r for all r , then it is a competitive allocation.*

We give an outline of the proof. To simplify things we consider the two consumer

case. The proof can be split into two parts. First we construct an analytical tool that will tell us whether a feasible allocation is competitive. Then using this tool we will show that a core allocation for all replica economies is a competitive allocation.

Step 1: Let (x_1, x_2) be a feasible allocation. For each consumer $i = 1, 2$ let $\Gamma_i = \{z \in \mathbb{R}^\ell : z + \omega_i \succ_i x_i\}$. This is the set of net trades z that result in a commodity bundle preferred to x_i . Let Γ be the convex hull³ of the union of these sets. This situation is depicted in Figure C.1.




Figure C.1: *Forming the set Γ*

It turns out that an allocation (x_1, x_2) is competitive if and only if Γ and the origin are disjoint.

To see this suppose (x_1, x_2) is a competitive allocation. Then there is an equilibrium price system p such that every commodity bundle preferred to x_i by the i 'th consumer is unaffordable. This means Γ_i , and hence Γ , lie strictly above the hyperplane through the origin with normal p ⁴. Therefore Γ and the origin are

³Let S be a subset of \mathbb{R}^ℓ . The convex hull of S is the intersection of all the convex sets containing S .

⁴Let p be an element of \mathbb{R}^ℓ different from zero; the set $H = \{z \in \mathbb{R}^\ell : p \cdot z = 0\}$ is a hyperplane through the origin with normal p .

disjoint. This can be seen in Figure C.1.

Conversely, if Γ and the origin are disjoint, then we can obtain an equilibrium price system as a normal vector to a separating hyperplane through the origin. The intuition can again be seen in Figure C.1. The commodity bundles x_i will lie on the hyperplane because the allocation (x_1, x_2) is feasible. Furthermore, Γ_i will lie strictly above the hyperplane which implies that a commodity bundle preferred to x_i is not affordable with this price system. Hence (x_1, x_2) is a competitive allocation.

Step 2: The second step in the proof shows that if (x_1, x_2) is in the core of all replica economies, then the origin does not belong to Γ . This is shown by contradiction. Assume that the origin does belong to Γ . We will show this implies (x_1, x_2) is blocked. Because Γ is the convex hull of the union of sets Γ_i we can express the origin as a convex combination of points in these sets. I.e.

$$\alpha_1 z_1 + \alpha_2 z_2 = 0$$

for positive α_1 and α_2 summing to one where $z_i + \omega_i \succ_i x_i$.

Consider a sequence z_i^k defined by $[k\alpha_i/a_i^k]z_i$ where a_i^k is the smallest integer greater than or equal to $k\alpha_i$, and k is an integer. As k tends to infinity this sequence tends to z_i and so is eventually in Γ_i for all k larger than some integer K (Γ_i is open.) We then have

$$a_1^K z_1^K + a_2^K z_2^K = K(\alpha_1 z_1 + \alpha_2 z_2) = 0$$

and $z_i^K + \omega_i \succ_i x_i$ for $i = 1, 2$. This implies that the coalition consisting of a_i^K members of type $i = 1, 2$ each of whom get $z_i + \omega_i$ blocks (x_1, x_2) in the economy replicated $\max\{a_1^K, a_2^K\}$ times. A contradiction. So the origin and Γ are disjoint and thus (x_1, x_2) is a competitive allocation.

We now prove the result formally.

Theorem 7. *If $x = (x_1, \dots, x_m)$ is in the core of \mathcal{E}^r for all r , then it is a competitive allocation.*

Proof. Let Γ_i be the set of all commodity bundles z in the commodity space such that $z + \omega_i \succ_i x_i$ and let Γ be the convex hull of the union of each Γ_i . Suppose by contradiction that the origin belongs to Γ . By convexity of preferences each Γ_i is a convex set. Thus $\sum_{i=1}^m \alpha_i z_i = 0$ for some $\alpha_i \geq 0$, $\sum_{i=1}^m \alpha_i = 1$ and $z_i + \omega_i \succ_i x_i$. Let a_i^k be the smallest integer greater than or equal to $k\alpha_i$, where k is an integer, and let I be the set of i such that $\alpha_i > 0$. For each i in I define a sequence z_i^k by $[k\alpha_i / a_i^k]z_i$. This sequence converges to z_i as k tends to infinity. Since preferences are continuous the sequence is *eventually* in Γ_i . Let K be an integer such that k larger than K implies z_i^k belongs to Γ_i for each i in I . Consider the coalition of a_i^K members of each type i to each of whom we assign $z_i^K + \omega_i$. Notice

$$\sum_{i \in I} a_i^K (z_i^K + \omega_i) = K \sum_{i \in I} \alpha_i z_i + a_{iK} \sum_{i \in I} \omega_i = a_{iK} \sum_{i \in I} \omega_i.$$

Such a coalition blocks x in the economy replicated $\max\{a_i^K : i \in I\}$ times, contradicting the assumption that x is in the core for all r .

Thus the origin does not belong to Γ . By Minkowski's separating hyperplane theorem [18, p.25] there is a hyperplane through the origin with normal p such that $p \cdot z \geq 0$ for all z in Γ . By continuity of preferences $x_i - \omega_i$ belongs to the closure of Γ so that $p \cdot x_i \geq p \cdot \omega_i$. Since

$$\sum_{i=1}^m x_i = \sum_{i=1}^m \omega_i$$

$p \cdot x_i = p \cdot \omega_i$ for each i . If $y \succ_i x_i$, then $y - \omega_i$ belongs to Γ so that $p \cdot y \geq p \cdot \omega_i$. Suppose by contradiction that $p \cdot y = p \cdot \omega_i$. By continuity and monotonicity of preferences, and the assumption that ω_i is strictly positive in each of its coordinates, there exists y' such that $y' \succ_i x_i$ and $p \cdot y' < p \cdot \omega_i$. A contradiction. \square

C.1.3 A model of perfect competition

[5] describes an exchange economy in which there is an infinite number of consumers, each consumer owns a negligible fraction of the total resources of each commodity, and each consumer acts independently of other consumers. He captures these conditions by modelling the set of consumers as an atomless measure space. In Aumann's model each consumer is negligible and independent, thus it is a model of perfect competition.

We first formally define the economy and then make a few remarks about the implications and interpretation of the model. [30] is the appropriate reference for the measure theoretic ideas.

The economy

The model of an economy consists of a set of consumers each of whom possess a commodity bundle and hold a preference relation over the set of all commodity bundles. The space \mathbb{R}^ℓ is called the *commodity space*, where ℓ refers to the number of commodities. A *commodity bundle* is an element of the commodity space. A *preference relation*, denoted \succeq , is a complete preordering over the commodity space and induces a *strict preference relation* (denoted \succ) in the following way: If x and y are arbitrary commodity bundles, then

$$x \succ y \text{ if and only if } x \succeq y \text{ and not } y \succeq x.$$

The set of all such preference relations is \mathcal{P} . $x \succeq y$ is interpreted to mean "x is preferred or indifferent to y."

The set of consumers is denoted T . There is a positive and finite measure

space⁵ (T, \mathcal{T}, μ) over the set of consumers. An *allocation* is an integrable function from T into the non-negative orthant of the commodity space. There is a fixed allocation, denoted ω , called the *initial allocation*; $\omega(t)$ is called the *initial endowment* of consumer t . An allocation x is said to be *feasible* if $\int x = \int \omega$. Each consumer has a preference relation defined by a *preference function* $\succeq: T \rightarrow \mathcal{P}$. We say that $\succeq: T \rightarrow \mathcal{P}$ is a *measurable preference function* if for arbitrary allocations x and y $\{t \in T : x(t) \succ_t y(t)\}$ is measurable.

Definition 10. An economy, denoted \mathcal{E} , is a triad $((T, \mathcal{T}, \mu), \omega, \succeq)$. It consists of a positive and finite measure space (T, \mathcal{T}, μ) , an initial allocation $\omega : T \rightarrow \mathbb{R}_+^\ell$, and a measurable preference function $\succeq: T \rightarrow \mathcal{P}$.

Let (T, \mathcal{T}, μ) be a measure space. An *atom* is a non-null element of the σ -algebra with the property that any subset also belonging to the σ -algebra is of equal or zero measure. That is, $A \in \mathcal{T}$ is an atom if for all $B \subseteq A$, where $B \in \mathcal{T}$, we have $\mu(B) = \mu(A)$ or otherwise $\mu(B) = 0$. A measure space with atoms is called *atomic*. A measure space without atoms is called *atomless*.

Core and competitive equilibrium

A *price system*, denoted p , is a vector in \mathbb{R}_+^ℓ .

Definition 11. A competitive equilibrium (p, x) is a price system p and an allocation x such that x is feasible and for μ -almost all t in T

1. $p \cdot x(t) \leq p \cdot \omega(t)$ and

⁵If μ is a measure on a σ -algebra \mathcal{T} , a set E in \mathcal{T} is said to have finite measure if $\mu(E) < \infty$. If the measure of every set E in \mathcal{T} is finite, then the measure μ is called finite. The measure μ is called positive if for all E in \mathcal{T} we have $\mu(E) \geq 0$. A positive and finite measure space is a measure space that has a positive and finite measure.

2. $y \succ_t x(t)$ implies $p \cdot y > p \cdot \omega(t)$.

If (p, x) is a competitive equilibrium, then p is called an *equilibrium price system* and x is called a *competitive allocation*.

The second equilibrium concept we consider is called the core. A *coalition* is any set S in \mathcal{T} with positive measure. We say that an allocation x is *blocked* by a coalition S if there exists an allocation y such that $\int_S y = \int_S \omega$, and $y(t) \succ_t x(t)$ for all t in S .

Definition 12. *The core is the set of feasible allocations not blocked by any coalition.*

Aumann's model of perfect competition

Aumann's model consists of an economy $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ such that (T, \mathcal{T}, μ) is an atomless measure space. We make the following remarks.

Remark 1. *Integration.* The main conceptual difference between Aumann's model and the Arrow-Debreu model is the use of integration. It is instructive to consider an example to see what the integral of an allocation means. In the Arrow-Debreu model the total allocation is obtained by summing the m commodity bundles; in Aumann's model we integrate.

The integral of an allocation is the *per capita* allocation: the allocation that each consumer obtains if we divide the allocation equally. Consider the following example. Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ be an economy with a single commodity. Suppose that the allocation x assigns half the economy S two units of the commodity and everyone else $T \setminus S$ nothing. The integral of this allocation is

$$\int x = 2\mu(S) + 0\mu(T \setminus S) = \mu(T).$$

Thus the per capita allocation is $\mu(T)$ units of the commodity. We can also define the mean allocation by dividing the per capita allocation by the measure of the

economy. In this case the mean allocation of the commodity is one. Notice that if the measure of the economy is one, then the per capita allocation and the mean allocation are the same.

Remark 2. *Core and Competitive Equilibrium.* The definition of a competitive equilibrium implies that the commodity bundle of any finite number of consumers in the economy can be changed arbitrarily – finite groups of consumers have measure zero and so have no effect. This may appear to be a weakening of the notion of a competitive allocation but it is not. The reason is that an equilibrium price system describes an infinite number of allocations – any allocation in which almost all of the consumers in the economy choose a maximal element in their budget set. This includes the allocation, induced by the equilibrium price system, in which all consumers in the economy choose a maximal element of their budget set.

For the core, the restriction to coalitions with positive measure may seem even more strange. A redistribution of resources in which a null set of consumers receive a commodity bundle inferior to their initial endowment may be in the core. But again, the core allocations need to be thought of as an equivalence class: an entire set of allocations in which the commodity bundles assigned to any null set of consumers in the economy can be changed arbitrarily. This includes the allocation where no null set of consumers can block.

Remark 3. *Lyapunov's Convexity Theorem.* An important tool used to study Aumann's model is Lyapunov's convexity theorem. This states that the range of an atomless measure is a closed convex subset of Euclidian space. In terms of the model the theorem says the following. Let S be a coalition and $0 < c < \mu(S)$. Then there is a coalition $A \subseteq S$ such that $\mu(A) = c$. This result is critical in proving that the set of competitive allocations coincides with the core.

Remark 4. *Implications of the Model.* At the beginning of this section we defined three features of Aumann's model: (1) there is an infinite number of consumers; (2) each consumer owns a negligible fraction of total resources; (3) each consumer acts independently of other consumers. We now describe how his model of an economy characterises these conditions.

(1) *There are an infinite number of consumers:* The set of consumers has positive measure. Since each consumer has zero measure there is an uncountable infinity of them.

(2) *Each consumer owns a negligible fraction of total resources:* This follows directly from the fact that each consumer has measure zero. The consequence is that an allocation's feasibility does not depend on what any individual consumer does.

(3) *Each consumer acts independently of other consumers:* This is implied by the model insofar that we can 'split' any coalition into smaller coalitions. The condition that consumers do not collude when interacting with the market is assumed rather than an endogenous feature of the model.

C.1.4 Aumann: Markets with a continuum of traders

Statement of the theorem

Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succ_t)$ be an economy where (T, \mathcal{T}, μ) is an atomless measure space and the following assumptions hold.

1. *Monotonicity:* let x and y be arbitrary commodity bundles. Then $x > y$ implies $x \succ_t y$;
2. *Continuity:* let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_t z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_t x\}$ are open;

3. $\omega(t)$ is strictly positive in each of its coordinates.

Under these conditions Aumann proves the following.

Theorem 8. *The core coincides with the set of competitive allocations.*

The last assumption that the initial endowment of each consumer is strictly positive in each of its coordinates is unnecessarily restrictive. We could have gotten away with assuming instead that each commodity is present in the market (i.e. $\int \omega$ is strictly positive in each of its coordinates.) We make this assumption to simplify the exposition of the proof. There is no loss here; the extra work that needs to be done under the less restrictive assumption is straightforward but a distraction from the proofs central idea.

A competitive allocation is in the core

Every competitive allocation belongs to the core. The argument for this result mimics its proof in the Arrow-Debreu model: if a coalition can block a competitive allocation, then at the equilibrium price system the total value of the coalition's commodity bundles exceeds the total value of their initial endowments.

Theorem 9. *A competitive allocation is in the core.*

Proof. Let x be a competitive allocation and suppose by contradiction that x does not belong to the core. Then there is a coalition S that blocks x with y . $y(t) \succ_t x(t)$ implies $p \cdot y(t) > p \cdot \omega(t)$ for each t in S . Hence, by the monotonicity and linearity of the integral

$$p \cdot \int_S y > p \cdot \int_S \omega,$$

contradicting $\int_S y = \int_S \omega$. □

A core allocation is a competitive allocation

The interesting part of Aumann's proof is in showing that a core allocation is a competitive allocation. The basic logic follows [19]. Aumann proves

Theorem 10. *A core allocation is a competitive allocation.*

We give an outline of the proof. Let x be a core allocation and Γ_t the set of net-trades for the t 'th consumer that result in a commodity bundle preferred to $x(t)$:

$$\Gamma_t = \{z \in \mathbb{R}_+^\ell : z + \omega(t) \succ_t x(t)\}.$$

Let Γ be the convex hull of the union of these sets taken over μ -almost all consumers, i.e. a set of consumers U such that $T \setminus U$ is a null set.

In precisely the same way as [19] it can be shown that a feasible allocation is a competitive allocation if and only if the origin does not belong to Γ .

Like in [19] we show that x is a competitive allocation by showing that the origin does not belong to Γ . Suppose by contradiction that the origin does belong to Γ . We will show this implies a coalition blocks the core allocation x . In [19] such a coalition is constructed by expressing the origin as a convex combination of points in sets Γ_i , where i refers to a particular type of consumer, and forming a coalition of sufficiently many consumers of these types by replication. This approach does not work for Aumann's model since (a) replicating the economy does not help us, and (b) a finite set of consumers is not a coalition (since it is a null set.)

Instead, Aumann expresses the origin as a convex combination of points

$$\sum_{i=1}^k \alpha_i z_i = 0$$

where each z_i is drawn from the set Γ_{t_i} with $t_i \in U$, and searches for non-null sets of consumers similar to t_i . Specifically he considers consumers that prefer $z_i + \omega(t)$

to their commodity bundle $x(t)$. I.e. $A(z_i) = \{t \in T : z_i + \omega(t) \succ_t x(t)\}$.

Assume for a moment that for each $i = 1, \dots, k$ $A(z_i)$ is non-null. This means that for each i there is a non-null set of consumers that prefers the net trade z_i to their own net trade. By Lyapunov's convexity theorem we can choose sets $S_i \subseteq A(z_i)$ such that $\mu(S_i) = \delta \alpha_i$ where δ is some positive real number chosen sufficiently small so that $\mu(A(z_i)) \leq \delta \alpha_i$ for each i . This situation is depicted in Figure C.2.




Figure C.2: *Forming a blocking coalition with Lyapunov's convexity theorem*

We could then form a coalition S equal to the union of the sets S_i and assign each consumer in S_i the commodity bundle $z_i + \omega(t)$. By construction, this commodity bundle is preferred to $x(t)$. Doing this is feasible for the coalition since

$$\int_S \sum_{i=1}^k (z_i + \omega) \chi_{S_i} = \delta \sum_{i=1}^k \alpha_i z_i + \int_S \omega = \int_S \omega.$$

Thus S blocks x contradicting that x is in the core.

So if it can somehow be shown that each set $A(z_i)$ is non-null the proof is complete. The way to do this is by careful choice of the set U . The general idea is to exclude consumers from U in such a way that $A(z_i)$ must be non-null. However we cannot exclude too many consumers from U otherwise this set will not satisfy

the property that $T \setminus U$ is a null set, or we will not be able to express the origin as a convex combination of commodity bundles from sets Γ_{t_i} where $t_i \in U$.

Two facts suggest a solution to the problem: (1) The rational numbers are dense in the real numbers, and (2) the rational numbers are countable.

The first fact suggests that we might be able to express the origin as a convex combination of commodity bundles with rational coordinates. That is

$$\sum_{i=1}^k \alpha_i r_i = 0$$

where $r_i \in \mathbb{Q}^\ell$, $r_i \in \Gamma_{t_i}$, $\alpha_i \geq 0$, and $\sum_{i=1}^k \alpha_i = 1$.

The second fact implies that the union of consumers $A(r)$ where $A(r)$ is a null set and r is a commodity bundle with rational coordinates is itself a null set. This follows because each null set is measurable, the set of measurable sets forms a σ -algebra, and the measure is countably additive. Thus setting

$$U = T \setminus \bigcup \{A(r) : r \in \mathbb{Q}^\ell \text{ and } \mu(A(r)) = 0\},$$

U is measurable and has the desired property that $T \setminus U$ is a null set. Suppose we do this.

Now pick any commodity bundle r with rational coordinates such that for some consumer t in U , r belongs to Γ_t ; that is $r + \omega(t) \succ_t x(t)$. By construction $A(r)$ is non-null. Since we can express the origin as the convex combination of commodity bundles with rational coordinates we can proceed as before to find a blocking coalition. Thus x is a competitive allocation.

We have not provided the details of how the origin is expressed as a convex combination of commodity bundles with rational coordinates. Doing this is a bit messy. The formal proof is given below.

Theorem 11. *A core allocation is a competitive allocation.*

Proof. Let x be a core allocation. For each commodity bundle z define $A(z)$ to be the set of consumers in T such that $z + \omega(t) \succ_t x(t)$ and define N to be the set of commodity bundles r with rational coordinates such that $A(r)$ is a null set. N is countable. Define the set of consumers U to be all consumers in T except those belonging to $A(r)$ where r is in N . U is measurable and its complement is a null set. Let Γ_t be the set of all commodity bundles z in the commodity space such that $z + \omega(t) \succ_t x(t)$ and let Γ be the convex hull of the union of each Γ_t , where the union is taken over the set U . Suppose by way of contradiction that the origin belongs to Γ . Then there exists $-y$ in Γ such that each component of y is positive. By the definition of Γ , $-y = \sum_{i=1}^k \gamma_i y_i$ for some $\gamma_i > 0$, $\sum_{i=1}^k \gamma_i = 1$ and $y_i + \omega(t_i) \succ_{t_i} x(t_i)$ where each t_i belongs to U . By continuity of preferences there exists rational numbers $\beta_i > 0$, $\sum_{i=1}^k \beta_i = 1$, and commodity bundles with rational coordinates r_i such that $-r = \sum_{i=1}^k \beta_i r_i$ where $r_i + \omega(t_i) \succ_{t_i} x(t_i)$ and each coordinate of r is positive. Choose an arbitrary consumer t_0 in U . By monotonicity there exists $\alpha > 0$ such that $\alpha r + \omega(t_0) \succ_{t_0} x(t_0)$. Set $r_0 = \alpha r$, $\alpha_0 = 1/(\alpha + 1)$, and $\alpha_i = \alpha \gamma_i / (\alpha + 1)$ for $i = 1, \dots, k$. Then $\alpha_i > 0$, $\sum_{i=0}^k \alpha_i = 1$, and $\sum_{i=0}^k \alpha_i r_i = 0$. By construction $A(r_i)$ is non-null for $i = 0, \dots, k$. For sufficiently small $\delta > 0$ each set $A(r_i)$ has a measure of at least $\delta \alpha_i$. By Lyapunov's convexity theorem there exists measurable subsets S_i of $A(r_i)$ with measure $\delta \alpha_i$. Let S be the union of sets S_i and define $y(t) = r_i + \omega(t)$ for each t in S_i . By construction $y(t) \succ_t x(t)$ for each t in S . Moreover,

$$\int_S y = \sum_{i=0}^k \delta \alpha_i r_i + \int_S \omega = \int_S \omega.$$

So coalition S blocks x with y contradicting that x belongs to the core.

Thus the origin does not belong to Γ . By Minkowski's separating hyperplane theorem [18, p.25] there is a hyperplane through the origin with normal p such that $p \cdot z \geq 0$ for z in Γ . By continuity of preferences $x(t) - \omega(t)$ belongs to the closure

of Γ so that $p \cdot x(t) \geq p \cdot \omega(t)$ for μ -almost all t . Since

$$\int x = \int \omega$$

$p \cdot x(t) = p \cdot \omega(t)$ for μ -almost all t .

If $y \succ_t x(t)$, then $y - \omega(t)$ belongs to Γ so that $p \cdot y \geq p \cdot \omega(t)$. Suppose by contradiction that $p \cdot y = p \cdot \omega(t)$ for a non-null set of consumers. By continuity and monotonicity of preferences, and the assumption that $\omega(t)$ is strictly positive in each of its coordinates, there exists y' such that $y' \succ_t x(t)$ and $p \cdot y' < p \cdot \omega(t)$ for a non-null set of consumers. A contradiction. \square

C.2 Imperfect Competition

This appendix studies imperfectly competitive economies. These are economies where individual consumers are non-negligible or groups of consumers act together as collusive entities. The mathematical model of this collusion is given by the notion of an *atom*. We first define the model of imperfect competition and discuss some conceptual issues that arise from the introduction of atoms. We then review two important papers: each establishes core equivalence in imperfectly competitive economies satisfying some restriction. The first paper by [50] proves core equivalence for an imperfectly competitive economy where atoms have the same preferences and initial endowment. The second paper by [24] proves core equivalence for an imperfectly competitive economy where each type of atom can be associated with a sufficiently large set of non-atomic consumers of their type. In particular Gabszewicz and Mertens paper explains why Aumann's core equivalence proof does not work for an imperfectly competitive economy.

C.2.1 Imperfect competition

We have defined a perfectly competitive economy to be one in which each consumer is negligible and behaves independently of other consumers. Thus an imperfectly competitive economy is one in which each consumer is non-negligible or where consumers do not behave independently of one another. We will argue that with infinitely many consumers the concept of a non-negligible consumer is nonsensical. Thus the way in which we introduce imperfect competition into the model of an economy is to allow groups of consumers to collude.

C.2.2 A model of imperfect competition

Aumann's model of an economy is a model of perfect competition because each consumer is negligible and independent. Both of these characteristics are captured by the measure space of consumers being atomless. An imperfectly competitive economy arises when one or more of these properties fails.

The model of imperfect competition we consider relaxes the assumption that consumers are independent. It introduces collusive groups that control a non-negligible fraction of the resources of the economy and thus can potentially manipulate prices. The model has three basic features: there are an infinite number of consumers, each consumer is negligible, and consumers are not independent in the sense that certain groups are assumed to interact cooperatively or competitively as a single entity.

This collusion is introduced into the model exogenously – we assume that certain groups of agents are bound together (like in a firm or trade union.) This is in contrast to the more natural approach of allowing collusion to form endogenously according to the mutual interests of groups of agents. Such an approach, while natural, is impossible.

Our model of perfect competition is an economy $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ where the measure space (T, \mathcal{T}, μ) is atomic. Recall that a measure space is atomic if the σ -algebra contains atoms and that an atom is a non-null element of the σ -algebra with the property that any subset of the atom belonging to the σ -algebra has equal or zero measure. The atom is what makes our model of an economy imperfectly competitive. We make the following remarks about the model.

Remark 5. *An atom represents a group of consumers:* We view an atom as a group of consumers collectively owning a non-negligible fraction of the total resources of the economy and bound together by common preferences. The size of the atom is given by its measure and this corresponds to the size of the group of consumers. Since the atom is non-null it is composed of an infinite number of consumers.

A commodity bundle for an atom represents a commodity bundle for each of its members. Thus the interpretation of a commodity bundle is different for an atom and a non-atomic consumer.

Given a price system the atom may sell the initial endowments of its members and purchase the most desirable element of its budget set according to the common preferences of its members. Thus the atom behaves as to maximise the welfare of its members much like a firm would for its shareholders. Formally the decision process of an atom is identical to one of its members. The only difference is that the atom has positive measure.

Nevertheless, it does not make sense to think of an atom as an individual consumer. If we pursue this interpretation, then what we are saying is that an individual consumer owns a non-negligible fraction of the resources of the economy. In general, with infinitely many consumers the quantity of each commodity is infinite so that a non-negligible fraction of resources involves infinite quantities of commodities. Besides it being awkward to assume that a single consumer owns

infinite quantities of certain commodities it becomes impossible to define reasonable preferences for them.

For example, suppose that faced with a price system a non-negligible consumer wishes to sell his initial endowment and purchase the most desirable commodity bundle affordable to him. If his preferences are defined over finite quantities of commodities, then there is no solution to this problem. For a solution to be well defined his preferences must be defined over commodity bundles that refer to infinite quantities of commodities. An individual consumer would not have such preferences.

Remark 6. *Preferences are μ -a.e. constant on an atom:* We now establish that the preference function is constant μ -almost everywhere on an atom. This follows from the preference function being measurable.

Proposition 6. *The preference function is μ -a.e. constant on an atom.*

Proof. To see this let \mathcal{E} be an economy and A an atom of the measure space (T, \mathcal{T}, μ) . Suppose by contradiction that the preference function is not μ -a.e constant on A . Then there exist commodity bundles x and y such that $C = \{t \in A : x \succ_t y\}$ is non-null and $A \setminus C$ is non-null. By the measurability of the preference function, since A is measurable, and the class of measurable sets \mathcal{T} is a σ -algebra we have that C and $A \setminus C$ are measurable. But this contradicts that A is an atom since

$$0 < \mu(C) < \mu(A) - \mu(A \setminus C) < \mu(A).$$

□

Remark 7. *Allocations are μ -a.e. constant on an atom:* This follows from the integrability of an allocation. If an allocation assigned different commodity bundles to

non-null subsets of an atom, then the allocation would not be measurable and hence not integrable.

C.2.3 Shitovitz: Oligopoly in markets with a continuum of traders

[50] proves core equivalence for an imperfectly competitive economy with at least two atoms having the same convex preferences and initial endowment. The basic approach is to reduce an imperfectly competitive economy to an equivalent perfectly competitive economy and to apply the core equivalence theorem of [5]. Shitovitz constructs a perfectly competitive economy by ‘splitting’ the atoms of his imperfectly competitive economy into atomless sets of consumers. He shows that core and competitive allocations are essentially the same for both economies. Having made the necessary assumptions he then applies Aumann’s core equivalence theorem which yields core equivalence in the imperfectly competitive economy.

Statement of the theorem

Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succ)$ be an economy where (T, \mathcal{T}, μ) is an atomic measure space with at least two atoms in which each atom has the same convex preferences and initial endowment. In order to apply the core equivalence result of [5] we assume

1. *Monotonicity*: let x and y be arbitrary commodity bundles. Then $x > y$ implies $x \succ_t y$;
2. *Continuity*: let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_t z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_t x\}$ are open; and
3. $\int_T \omega(t) >> 0$ (i.e. every commodity exists.)

Under these conditions it can be shown that

Theorem 12. *The core coincides with the set of competitive allocations.*

The assumption that atoms have convex preferences is crucial.

Proof of the theorem

We follow the proof given by [27]. Let T_0 be the atomless part of T . We assume T_0 is non-null so without loss of generality $T_0 = [0, 1]$. Denote the atomic part of T by $T_1 = T \setminus T_0$. Define an economy \mathcal{E}^* by $((T^*, \mathcal{T}^*, \mu^*), \omega^*, \succeq^*)$ where $T^* = [0, 1 + \mu(T_1)]$ and \mathcal{T}^* and μ^* are given by the direct sum of \mathcal{T} and μ restricted to T_0 and the Lebesgue atomless measure space over $[1, 1 + \mu(T_1)]$. Given an atom $A \subseteq T$ we define the ‘split’ atom $A^* \subseteq T^*$ to be an atomless set such that $\mu^*(A^*) = \mu(A)$. Denote the set of split atoms by T_1^* . Let $T_0^* = T^* \setminus T_1^*$.

We define the initial allocation and the preference function of \mathcal{E}^* by

$$(\omega^*(t), \succeq_t^*) = \begin{cases} (\omega(A), \succeq_A) & \text{if } t \in A^*, \\ (\omega(t), \succeq_t) & \text{if } t \in T_0^*. \end{cases}$$

In words, each consumer in \mathcal{E} that belongs to an atom is obliged to work in unison with the other members of the atom. \mathcal{E}^* is obtained by removing this restriction. The idea is summarised in Figure C.3. On the left is an imperfectly competitive economy with two atoms, on the right is the economy after the atoms have been split.

To compare core and competitive allocations between economies we define a mapping from the set of allocations of one economy to the set of allocations of the other. Given an allocation x for \mathcal{E} an allocation x^* for \mathcal{E}^* is given by distributing the commodity bundle of each atom equally among its members. Thus for each t in T^*




Figure C.3: From imperfect to perfect competition

$$x^*(t) = \begin{cases} x(t) & \text{if } t \in T_0, \\ x(t) & \text{if } t \in A^*, A \subseteq T_1 \text{ and } A \in \mathcal{T}. \end{cases}$$

In the other direction, given an allocation x^* for \mathcal{E}^* the allocation x for \mathcal{E} obtains by averaging the commodity bundles of consumers in A^* . Thus for each t in T

$$x(t) = \begin{cases} x^*(t) & \text{if } t \in T_0, \\ \frac{1}{\mu^*(A^*)} \int_{A^*} x^* & \text{if } t \in A \subseteq T_1. \end{cases}$$

It is clear that an allocation for one economy is feasible if and only if it is feasible for the other economy.

We first prove some preliminary results needed for the proof of Theorem 12. The convexity of each atom's preferences implies that if an allocation assigns each consumer in the split atom a commodity bundle preferred to the commodity bundle x , then the mean of these commodity bundles is also preferred to x .

Lemma 7. Let \succeq_{A^*} be a convex preference relation. If f is an allocation such that $f(t) \succeq_{A^*} x$ for all t in A^* , then $\frac{1}{\mu^*(A^*)} \int_{A^*} f \succeq_{A^*} x$.

Proof. If f is a simple function, then the result follows by induction using the

convexity of preferences. Otherwise, let $\{f_n\}$ be a sequence of measurable functions on A^* such that $0 \leq f_1 \leq f_2 \leq \dots$, $f_n(t) \in f(A^*)$ for all t in A^* , and $\lim f_n = f$ in measure. The sequence $\{\frac{1}{\mu^*(A^*)} \int_{A^*} f_n\}$ belongs to the set $C = \{z \in \mathbb{R}^\ell : z \succeq_{A^*} x\}$ which is closed by continuity. By the Monotone convergence theorem

$$\lim \frac{1}{\mu^*(A^*)} \int_{A^*} f_n = \frac{1}{\mu^*(A^*)} \int_{A^*} f$$

and so $\frac{1}{\mu^*(A^*)} \int_{A^*} f$ belongs to C . □

We state the next result without proof. It is due to Karl Vind ([55].) Fortunately the result is quite intuitive. It states that if an allocation is blocked by some coalition in a perfectly competitive economy (i.e. an economy with an atomless set of consumers), then it can be blocked by a coalition of any measure less than the measure of the whole set of consumers.

Lemma 8. *Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ be an economy where (T, \mathcal{T}, μ) is an atomless measure space. If an allocation x is blocked by a coalition S , then for any c , $0 < c < \mu(T)$, there exists a coalition E with $\mu(E) = c$ that blocks x .*

We can now prove Theorem 12. We do so in two parts. First, we show that an allocation for \mathcal{E} is competitive if and only if the associated allocation is competitive for \mathcal{E}^* . Then we show that an allocation for \mathcal{E} is in the core if and only if the associated allocation is in the core for \mathcal{E}^* .

Lemma 9. *(p, x) is a competitive equilibrium for \mathcal{E} if and only if (p, x^*) is a competitive equilibrium for \mathcal{E}^* .*

Proof. It is clear that an allocation for one economy is feasible if and only if it is feasible for the other economy.

Let (p, x) be a competitive equilibrium for \mathcal{E} . x^* assigns each consumer in \mathcal{E}^* the commodity bundle he receives in \mathcal{E} . Thus (p, x^*) is a competitive equilibrium for \mathcal{E}^* .

Conversely, let (p, x^*) be a competitive equilibrium for \mathcal{E}^* . Recall that x averages the commodity bundles of consumers belonging to the split atom A^* . The only difficulty in this case is that the members of the split atom may receive different commodity bundles. However, since each consumer belonging to the split atom has convex preferences Lemma 7 implies that the average is preferred or indifferent to the commodity bundle under x^* . By the definition of competitive equilibrium

$$p \cdot x^*(t) \leq p \cdot \omega^*(t) \text{ and } y \succ_t^* x^*(t) \text{ implies } p \cdot y > p \cdot \omega^*(t).$$

for μ -almost all consumers in T^* . Since

$$p \cdot \frac{1}{\mu(A_1^*)} \int_{A_1^*} x^* \leq p \cdot \omega(t)$$

we have $p \cdot x(A) \leq p \cdot \omega(A)$ for each atom in T_1 and $p \cdot x(t) \leq p \cdot \omega(t)$ for μ -almost all t in T_0 .

Let $y \succ_t x(t)$ for some $t \in T$. Then since $\succeq_t = \succeq_t^*$ and $\omega(t) = \omega^*(t)$ we have

$$p \cdot y > p \cdot \omega(t)$$

for μ -almost all t in T . Thus (p, x) is a competitive equilibrium for \mathcal{E} . □

This completes the first part of the proof. The next part of the proof is more difficult.

Lemma 10. *x is a core allocation for \mathcal{E} if and only if x^* is a core allocation for \mathcal{E}^* .*

Proof. Let x^* be a core allocation for \mathcal{E}^* . Suppose by contradiction that x is not a

core allocation for \mathcal{E} . Then there is a coalition S that blocks x with y . But y^* is feasible and preferred to x^* by each member of S^* . Thus S^* blocks x^* with y^* . A contradiction.

Conversely, assume that x is a core allocation for \mathcal{E} . Suppose by contradiction that x^* is not a core allocation for \mathcal{E}^* . Then there exists a coalition S^* that blocks x^* with y^* . We will show this implies that a coalition R blocks x .

Since each atom has the same preferences there exists an atom \bar{A} that receives the least desired commodity bundle. That is,

$$x(A) \succeq_A x(\bar{A})$$

for all atoms $A \subseteq T_1$. Let $\mu(\bar{A}) = \alpha$.

Since there are at least two atoms we have $\mu^*(T^*) > 1 + \alpha$. By Lemma 8 ([55]) we can choose the measure of our blocking coalition to be $1 + \alpha$. So assume $\mu^*(S^*) = 1 + \alpha$.

Since T_0^* has measure one the measure of the coalition intersected with T_1^* must be at least α by subadditivity. That is, $\mu^*(S^* \cap T_1^*) \geq \alpha$. We will show using Lyapunov's convexity theorem that there exists a coalition $R^* \subseteq S^*$ such that $\mu(R^* \cap T_1^*) = \alpha$ and that this coalition blocks x^* . To see this define the vector measure ν on T^* by

$$\nu(E) = \left(\int_E (y^* - \omega^*), \mu^*(E \cap T_1^*) \right)$$

for any $E \in \mathcal{T}^*$. Since $(T^*, \mathcal{T}^*, \mu^*)$ is an atomless measure space Lyapunov's theorem implies there is a coalition $R^* \subseteq S^*$ such that

$$\nu(R^*) = \frac{\alpha}{\mu^*(S^* \cap T_1^*)} \nu(S^*).$$

Therefore

$$\begin{aligned}
v(R^*) &= \left(\int_{R^*} (y^* - \omega^*), \mu^*(R^* \cap T_1^*) \right) \\
&= \left(\frac{\alpha}{\mu^*(S^* \cap T_1^*)} \int_{S^*} (y^* - \omega^*), \frac{\alpha \mu^*(S^* \cap T_1^*)}{\mu^*(S^* \cap T_1^*)} \right) \\
&= (0, \alpha)
\end{aligned}$$

which implies that y^* is feasible for R^* and $\mu(R^* \cap T_1^*) = \alpha$.

To finish the proof we define an allocation for $(R \cap T_0) \cup \bar{A}$. Let

$$y(t) = \begin{cases} y^*(t) & \text{if } t \in R \cap T_0, \\ \frac{1}{\alpha} \int_{R^*} y^* & \text{if } t = \bar{A}. \end{cases}$$

y is preferred to x for each member of $R \cap T_0$. By Lemma 7 y is preferred to x by \bar{A} .

Thus $(R \cap T_0) \cup \bar{A}$ blocks x with y . A contradiction. \square

By [5] the core of \mathcal{E}^* coincides with the set of competitive allocations. By Lemma 9 and Lemma 10 the core of \mathcal{E} coincides with the set of competitive allocations. This completes the proof.

C.2.4 Gabszewicz and Mertens: Atoms that aren't "too" big

[24] prove core equivalence for imperfectly competitive economies where each atom can be associated with a sufficiently large set of atomless consumers with the same preferences and initial endowment as the atom. Their theorem is a direct extension of [5]. The key to their result is showing that core allocations have an equal treatment property. This means that a core allocation assigns to consumers with the same preferences and initial endowment commodity bundles in the same indifference set.

Formally, let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ be an economy where (T, \mathcal{T}, μ) is a separable⁶ atomic measure space and each atom has convex preferences. By separability the set of atoms can be partitioned into at most countably many sets $\{T_{a,i} : i \geq 1\}$ where each $T_{a,i}$ contains atoms with the same preferences and initial endowment. For each $T_{a,i}$ define T_i to be the set of consumers with the same preferences and initial endowment as the atoms in $T_{a,i}$. Notice that $T_{a,i} \subseteq T_i$. Define the set of consumers not in any T_i as T_0 . $T_0 = T \setminus \bigcup\{T_i : i \geq 1\}$ is atomless.

So that they can apply the core equivalence result of [5] Gabszewicz and Mertens assume \mathcal{E} satisfies

1. *Monotonicity*: let x and y be arbitrary commodity bundles. Then $x > y$ implies $x \succ_t y$;
2. *Continuity*: let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_t z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_t x\}$ are open; and
3. $\int_T \omega(t) >> 0$ (i.e. every commodity exists.)

Under these conditions they prove:

Theorem 13. *If $\sum_{i \geq 1} \frac{\mu(T_{a,i})}{\mu(T_i)} \leq 1$ (holding strictly in the case of a single atom), then the core coincides with the set of competitive allocations.*

Why Aumann's proof does not work for imperfectly competitive economies

Recall Aumann's core equivalence proof. Aumann applies the argument of [19] to a perfectly competitive economy: he considers a core allocation x and then assumes by contradiction that x is not competitive, or equivalently that $0 \in \Gamma$ where

⁶A measure space (T, \mathcal{T}, μ) is called separable if its σ -algebra \mathcal{T} can be generated by a countable collection of sets.

$$\Gamma = co \bigcup_{t \in U} \{\Gamma_t : t \in U\} = co \bigcup_{t \in U} \{z : z + \omega(t) \succ_t x(t)\}$$

where U is a subset of consumers such that its complement in T is a null set.

He then proceeds to show that one can construct a blocking coalition to x . He finds commodity bundles r_1, \dots, r_k and weights $\alpha_1, \dots, \alpha_k$ such that

$$0 = \sum_{i=1}^k \alpha_i r_i$$

where each r_i has the property that $A(r_i) = \{t \in T : r_i + \omega(t) \succ_t x(t)\}$ is non-null. The idea is to select, by Lyapunov's theorem, a set of consumers S_i from each $A(r_i)$ with measure $\delta \alpha_i$, (δ being chosen sufficiently small.) The assignment of $r_i + \omega(t)$ to each member of S_i is feasible because

$$\sum_{i=1}^k \int_{S_i} (r_i + \omega(t)) = \delta \sum_{i=1}^k \alpha_i r_i + \sum_{i=1}^k \int_{S_i} \omega(t).$$

Thus x is blocked by the coalition made up of the union of each S_i .

Importantly, notice the assumption that (T, \mathcal{T}, μ) is an atomless measure space is used to apply Lyapunov's convexity theorem to each $A(r_i)$. The reason why Aumann's proof does not work in imperfectly competitive economies is that we can only apply Lyapunov's theorem to $A(r_i)$ if it contains a non-null set of atomless consumers. However it may be the case that $A(r_i)$ contains only atoms. In this case we cannot apply Lyapunov's theorem and so we cannot construct a blocking coalition. The goal of Gabszewicz and Mertens paper is to provide a sufficient condition such that the set $A(r_i)$ contains a non-null set of atomless consumers.

Assume for a moment that for each atom there exists a non-null set of atomless consumers with the same preferences and initial endowment as the atom. That is $\mu(T_i) > \mu(T_{a,i})$ for each $i \geq 1$. Moreover, assume that every core allocation assigns to each consumer in T_i a commodity bundle in the same indifference set – this is

the equal treatment property. Now consider the set of consumers

$$A(r_i) = \{t \in T : r_i + \omega(t) \succ_t x(t)\}.$$

If this set contains an atom (i.e. $r_i + \omega(A) \succ_A x(A)$ for some atom A), then it contains all consumers with the same preferences and initial endowment as the atom. Therefore $A(r_i)$ contains a non-null set of atomless consumers. Lyapunov's theorem can be applied and Aumann's proof works.

The hypothesis of Theorem 13 is

$$\frac{\sum_{i \geq 1} \mu(T_{a,i})}{\mu(T_i)} \leq 1$$

(with strict inequality in the case of a single atom.) Thus the first requirement, that $\mu(T_i) > \mu(T_{a,i})$, is satisfied. In the next section we establish that this hypothesis implies the desired equal treatment property. Once this has been achieved the proof is complete.

Idea of the proof

The goal is to establish that core allocations assign consumers of the same type commodity bundles in the same indifference set. The argument of the proof is similar to the proof of equal treatment in [19]. Assuming that the equal treatment property is false for some core allocation we will construct a blocking coalition consisting of a non-null set C_i of consumers from each type T_i . We choose C_i to contain those consumers that receive less desirable commodity bundles than other members of their type. We then assign each member of C_i the average of their type's commodity bundle. In order for this assignment to be feasible we must be able to choose the C_i to have equal measure. This is the main difficulty of the proof.

By Theorem 3.2.1 we may assume without loss of generality that the measure

of each T_i $i \geq 1$ is one. Consider each type T_i , $i \geq 1$, ordered on the unit interval according to the desirability of their commodity bundle (the consumers that receive the less desirable commodity bundles lying closer to the origin.) The ability to choose non-null sets C_i from each type so that all have equal measure is determined by the position of the atoms on the unit intervals (see Figure C.4.) In Figure C.4 the ability to choose such sets corresponds to whether we can find an $\alpha \in (0, 1)$ such that α does not intersect an atom. In the first panel of Figure C.4 we cannot do this because whichever α we choose it will always intersect one of the atoms. In the second panel, we can.

Figure C.4 refers to an economy with two types, T_1 and T_2 , and two atoms, A_1 and A_2 with $\mu(A_1) = \mu(A_2) = \frac{2}{3}$. Depending on the core allocation, an $\alpha \in (0, 1)$ that does not cut the atoms may or may not exist because if we line the atoms head to toe, their measure (here indicated by their length) spans the unit interval. These two cases are illustrated in the figure.




Figure C.4: *Forming a blocking coalition*

However if the sum of the measure of the atoms in the economy is less than or equal to one, then we will always be able to find an $\alpha \in (0, 1)$ that does not intersect the atoms. This is precisely the hypothesis of Theorem 13.

The proof of the theorem formalises this idea.

Proof of the theorem

We follow the proof of [42]. Let $\mathcal{E} = ((T, \mathcal{T}, \mu), \omega, \succeq)$ be an economy where (T, \mathcal{T}, μ) is a separable atomic measure space. Let each consumer belonging to T_i have the same utility function $u_i : \mathbb{R}_+^\ell \rightarrow \mathbb{R}$. Suppose that each atom has convex preferences. Assume \mathcal{E} satisfies

1. *Monotonicity*: let x and y be arbitrary commodity bundles. Then $x > y$ implies $x \succ_t y$;
2. *Continuity*: let x be an arbitrary commodity bundle. The sets $\{z \in \mathbb{R}_+^\ell : x \succ_t z\}$ and $\{z \in \mathbb{R}_+^\ell : z \succ_t x\}$ are open; and
3. $\int_T \omega(t) >> 0$ (i.e. every commodity exists.)

Lemma 11. *If $\sum_{i \geq 1} \frac{\mu(T_{a,i})}{\mu(T_i)} \leq 1$ (holding strictly in the case of a single atom) and x is a core allocation, then $u_i(x(t))$ is μ -a.e. constant on T_i for each $i \geq 1$.*

Proof. Let x belong to the core and suppose by contradiction that $u_i(x(t))$ is not μ -a.e. constant on all T_i . For each $i \geq 1$ define a measurable function $h_i : T_i \rightarrow [0, 1]$ such that $h_i(t) \neq h_i(s)$ for all $t \neq s$. For each $i \geq 1$ define the function $F_i : T_i \rightarrow [0, 1]$ by

$$F_i(t) = \mu(\{s \in T_i : u_i(t) > u_i(s), \text{ or } [u_i(t) = u_i(s) \text{ and } h_i(t) > h_i(s)]\}).$$

F_i is measurable by the measurability of u_i and h_i . For each atom $A \in T_{a,i}$ define the interval $J_A = (F_i(A), F_i(A) + \mu(A))$. Consider each T_i as ordered by F_i on the unit interval. By hypothesis $\sum_{i \geq 1} \mu(T_{a,i}) \leq 1$. Thus there exists $\alpha \in (0, 1)$ such that $\alpha \notin J_A$ for all atoms $A \subseteq T$. Define the sets $C_i = F_i^{-1}(\alpha)$. Clearly $\mu(C_i) = \alpha$

for all $i \geq 1$. By Lyapunov's convexity theorem choose a subset C_0 of T_0 such that $\int_{C_0}(x - \omega) = \alpha \int_{T_0}(x - \omega)$. Let $C = \cup\{C_i : i \geq 1\}$. Consider the following assignment for C :

$$y(t) = \begin{cases} x(t) & \text{if } t \in C_0, \\ \int_{T_i} x & \text{if } t \in C_i, i \geq 1. \end{cases}$$

By convexity of the atoms preferences and Lemma 7 $u_i(y(t)) \geq u_i(x(t))$ for all $i \geq 1$ and for at least one $i' \geq 1$ we have $u_{i'}(y(t)) > u_{i'}(x(t))$ for all $t \in C_{i'}$. y is feasible for C since

$$\begin{aligned} \int_C (y - \omega) &= \int_{C_0} (x - \omega) + \sum_{i \geq 1} \int_{C_i} \left(\left(\int_{T_i} x \right) - \omega \right) \\ &= \alpha \int_{T_0} (x - \omega) + \sum_{i \geq 1} \left[\mu(C_i) \left(\int_{T_i} x \right) - \int_{C_i} \omega \right] \\ &= \alpha \int_{T_0} (x - \omega) + \alpha \sum_{i \geq 1} \int_{T_i} (x - \omega) \\ &= \alpha \int_T (x - \omega). \end{aligned}$$

By monotonicity there exists an allocation \bar{y} for C such that $u_i(\bar{y}(t)) > u_i(x(t))$ for all $t \in C$ and $i \geq 0$ (simply redistribute a small quantity of the commodity bundle of consumers in $C_{i'}$ to the other members of C .) Thus C blocks x with \bar{y} . A contradiction. \square