Three Essays: Information in Decentralised Markets, Committees, and College Admissions



Dağhan Carlos Edip Akkar

Nuffield College

University of Oxford

August 2025

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Copyright © 2025 by Dağhan Carlos Edip Akkar

All Rights Reserved

This thesis was typeset using the Oxford Transfer of Status Template by Joseph Rowell and Frank Fu, available on Overleaf under a Creative Commons CC BY 4.0 license.

Acknowledgements

I am deeply grateful to my supervisors Ian Jewitt, Margaret Meyer, and Daniel Quigley.

Ian was an exceptionally generous and patient mentor. He inspired me to be a better scholar, through both his guidance and his example.

Meg contributed to my scholarly and personal development from the day she met me. Every aspect of my work benefited from her deep insights and generous feedback.

Dan became not just an exceptional supervisor, but also a friend. His mentorship, both academic and personal, was integral to my DPhil.

I am also deeply grateful to Ludvig Sinander for his mentorship and encouragement. He contributed immensely to my scholarly development, and served as a role model.

I warmly thank Manos Perdikakis for his friendship and companionship throughout graduate school. Every conversation with him renewed my appreciation for economics.

I thank the Clarendon Fund, Nuffield College, and the Department of Economics at Oxford for funding my DPhil. I am particularly grateful to Nuffield College for giving me an academic home and a welcoming community.

Finally, I thank my close friends, my partner Jasmin, and my family. My friends' support saw me through the many challenges of the DPhil. Jasmin was a tireless coach, cheerleader, and support beam for me. Without my family's help, my DPhil would not have been possible.

I benefited from comments and feedback provided by César Barilla, Inés Moreno de Barreda, Mikhail Drugov, Péter Eső, Stephan Lauermann, Paul Klemperer, Ludmila Matysková, Paula Onuchic, Manos Perdikakis, Daniel Rappoport, Marzena Rostek, Clara Schreiner, Ludvig Sinander, Alex Teytelboym, Andy Zapechelnyuk, seminar audiences at the Oxford Student Theory Workshop, NASMES 2024, EEA-ESEM 2024, the 2024 Paris Transatlantic Theory Workshop, and speakers at the Nuffield Economic Theory Seminar.

Sevgili ailem Mariela, Sinan, ve Defne'ye.

Table of Contents

Acknowledgements **i**

Abstract **1**

Introduction **3**

1 The (Mis)use of Information in Decentralised Markets **7**

1.1 Introduction 8

1.2 Model 15

1.3 Buyers' Beliefs, Equilibria, and Total Surplus 17

1.3.1 Buyers' Beliefs and Equilibria 17

1.3.2 Total Surplus 21

1.4 Efficiency with More Buyers 23

1.5 Efficiency with Better-Informed Buyers 26

1.5.1 Binary Experiments 27

1.5.2 Finite Experiments 33

1.6 Maximising Surplus by Coarsening Information 36

1.7 Ultimatum Price Offers by Buyers 40

1.8 Supplemental Appendix 44

1.9 Proof Appendix 46

1.9.1 Useful Definitions and Notation 46

1.9.2 Omitted Results 47

1.9.3 Omitted Proofs 52

References 82

2 Optimally Dictatorial Committees **86**

2.1 Introduction . . . 87

2.2 The Model . . . 92

2.3 Dictatorships . . . 96

2.4 A Dominant Dictatorship . . . 99

2.5 Discussion . . . 105

References . . . 110

3 Collusive College Admissions **113**

3.1 Introduction . . . 114

3.2 The Model . . . 119

3.3 Sequential Applications . . . 121

3.4 Simultaneous Applications . . . 124

3.4.1 Equilibrium Application Strategies and Matchings . . . 124

3.4.1.1 High Difference Between Investments: $(1-p)\times\alpha>\beta$ . . . 124

3.4.1.2 Low Difference Between Investments: $(1-p)\times\alpha<\beta$ . . . 126

3.4.2 Investments and Payoffs in Simultaneous Applications . . . 128

3.5 Equilibrium when Colleges Set Application Dates . . . 131

3.6 Conclusion . . . 132

3.7 Omitted Results . . . 134

3.8 Omitted Proofs . . . 140

References . . . 144

Conclusion **146**

List of Figures

1.5.1 Theorem 1.2 illustrated 28
1.5.2 Blackwell Improvements of a Binary Signal 29
1.5.3 A Local Mean Preserving Spread 34
1.8.a Total surplus 44
1.8.b Probability that buyer accepts upon $s = 0.8$ 44
1.8.c Probability that some buyer trades when $\theta = H$ 45
1.8.d Probability that some buyer trades when $\theta = L$ 45
1.8.e Probability that $\theta = H$ given some buyer trades 45
1.8.f Probability that $\theta = H$ given no buyer trades 45

3.3.1 The Positive Assortative Matching $\mu^s(a, \kappa)$ 122
3.3.2 Colleges' Expected Payoffs under Sequential Applications 123
3.4.1 The Equilibrium Matching in Simultaneous Adm.: $(1-p) \times \alpha > \beta$ 126
3.4.2 The Equilibrium Matching in Simultaneous Adm.: $\alpha > \beta > (1-p) \times \alpha$ . . . 127

List of Tables

1.1 Running Example: Comparing Equilibrium Strategies $\check{\sigma}$ and $\hat{\sigma}$ 19

1.2 Running Example: Comparing Surplus Across Equilibria 23

Abstract

This PhD dissertation presents three papers in the field of information economics.

In Chapter 1, I study the link between information and efficiency in a decentralised market. The market has one seller offering an asset of common value to prospective buyers. Each buyer has a private signal about the common value of the asset. I ask whether the market becomes allocatively more efficient with (i) more buyers, (ii) better-informed buyers. Both increase the information available about buyers' common value, but also the adverse selection each buyer faces. With more buyers, trade surplus eventually increases and converges to the full-information upper bound if and only if the likelihood ratios of buyers' signals are unbounded from above. Otherwise, it eventually decreases and converges to the no-information lower bound. With better information about trades buyers would have accepted, trade surplus increases. With better information about trades they would have rejected, trade surplus decreases—unless adverse selection is irrelevant. For binary signals, a sharper characterisation emerges: stronger good news increase total surplus, but stronger bad news eventually decrease it.

In Chapter 2, I study the optimal voting mechanism for a committee that must decide whether to enact or block a policy of unknown benefit. Information can come both from committee members who can acquire it at cost, and a strategic lobbyist who wishes the policy to be enacted. I show that the *dictatorship of the most-demanding member* is a *dominant* voting mechanism: any other voting mechanism is (i) less likely to enact a good policy, (ii) more likely to enact a bad policy, and (iii) burdens each member with a greater cost of acquiring information.

In Chapter 3, I study how caps on students' applications influence college admissions outcomes. I find that such caps may benefit both high- and low- ranked colleges. A

low-ranked college benefits from an *allocative effect* whereby it enrols better students. A high-ranked college benefits from a *competitive effect* which decreases the competition for the best students. This incentivises colleges to coordinate their application policies and prevent students from applying to both. I discuss how this helps explain the global persistence of college admission systems that generate unstable matchings.

Introduction

This PhD dissertation presents three stand-alone papers in the field of information economics.

In Chapter 1, I study the relationship between efficiency and information in a decentralised market. In such markets, trades occur through *bilateral* and *private* negotiations: each buyer meets the seller one-on-one, and has little knowledge of his previous failed negotiations. Among the most prominent examples are OTC markets for securities of unknown quality, venture capital markets for startups of unknown viability, and credit markets for borrowers of uncertain creditworthiness.

I ask: does a decentralised market become more efficient when buyers have more information about their common value for the asset being traded; i.e., the asset's *quality*? Specifically, does it become more efficient when:

1. there are more buyers, each with a private signal about the asset's quality?

2. each buyer receives a more informative signal about the asset's quality?

Both changes increase the information available about the asset's quality; buyers rely on this information to estimate the gains from trade. However, these changes also increase the *adverse selection* each buyer is exposed to: in a market with more- or better-informed buyers, each buyer faces a greater risk of accepting a trade which, were she aware of other buyers' negative assessments of the asset, she would deem unfavourable.

I show that the *kind* of information available in the market determines how this trade-off is resolved. Whether more buyers make the market more efficient depends on the signals buyers may receive about the asset's quality. If buyers may receive a signal which unambiguously reveals that the asset has high quality, a market with more buyers eventually becomes more efficient; market outcomes converge to full-information outcomes as the

number of buyers grows. Otherwise, a market with more buyers eventually becomes less efficient; total surplus in the market ultimately reaches the no-information benchmark.

Similarly, whether better-informed buyers make the market more efficient depends on the kind of improvement in buyers' information. *Negative overrides*—additional information about trades buyers would have accepted—always increase efficiency. On the other hand, *positive overrides*—additional information about trades buyers would have rejected—decrease efficiency; unless *adverse selection is irrelevant* in an appropriately defined sense.

Both my research question and my findings contribute to the literature on information aggregation in markets. Chapter 1 elaborates on this contribution and discusses the implications for the structure and regulation of financial markets.

In Chapter 2, I study the optimal voting mechanism for a committee which must decide whether to enact or block a policy of unknown benefit to society. Information about the policy can be acquired by committee members at personal cost, or provided by a strategic lobbyist who wishes the policy to be enacted. I call a voting mechanism *dominant* if it secures:

1. a higher chance of enacting a good policy,

2. a lower chance of enacting a bad policy, and

3. a lower expected cost of information acquisition for each member

than any other voting mechanism.

I show that a dominant voting mechanism exists by characterising one: *the dictatorship of the most-demanding member*. Under this mechanism, only one member's vote influences the outcome. That member, were she deciding without a committee or lobbyist to interfere, would obtain the most certainty that the policy is good before enacting it. I call her *the most-demanding member*.

This paper contributes to a vast interdisciplinary literature on optimal committee design. This literature studies how a committee does, or can be designed to, overcome four problems:

1. it must incentivise external stakeholders to provide information about its decision,

2. it must incentivise its members to acquire such information at their own cost,

3. it must minimise the cost its members bear as they acquire such information,

4. its final decision must efficiently aggregate the information it has accumulated.

The existing literature addresses these problems in isolation, focusing only on one or two of them at a time. Troublingly, the objectives each of these problems present are often in conflict. Therefore, papers which focus on different subsets of these problems offer conflicting guidance. This makes it difficult to draw policy insights from the existing literature, as policymaking committees often face all four of these problems simultaneously.

This paper is the first to study these four problems simultaneously. Despite the complexity of the problem, it identifies a simple voting mechanism and shows that it dominates every other voting mechanism the committee may design.

In Chapter 3, I study how capping the number of colleges students can apply to influences college admissions outcomes. My model has two crucial features. First, students' preferences towards colleges are endogenous: colleges compete to attract students by undertaking costly investments. Second, there is uncertainty about colleges' capacities; so, students are uncertain whether they would be admitted to a college if they applied. This creates a trade-off: when they are limited to a single application, students can either submit a risky application to the high-ranked college, or a safe application to the low-ranked college.

I find that two effects emerge when students are restricted to one application instead of being allowed to apply to both colleges. The first is an *allocative effect*: the low-ranked college attracts high-ability students who would have been admitted to the high-ranked college, but were unwilling to risk applying there. The low-ranked college benefits from this effect. As a consequence, the high-ranked college enrols worse students. Nonetheless, the high-ranked college benefits from the *competitive effect* this generates: since the low-ranked college receives a better share of students, it becomes less willing to compete for the very best students.

The allocative effect is consistent with existing findings in the college admissions literature. However, existing work concludes that this effect harms high-ranked colleges by causing them to lose desirable applicants. This conclusion stems from models that treat

students' preferences towards colleges as exogenous. In reality, however, colleges expend significant resources to maintain or improve their appeal among students. Motivated by this, my model endogenises students' preferences: colleges compete for better applicants by engaging in costly investments. This approach reveals that high-ranked colleges also benefit from application caps. This occurs through a competitive effect that significantly reduces the competition they face for the very best students. My finding helps explain the persistence of admission caps and costs in college admission systems around the world.

I conclude by brief discussing promising avenues for future research.

Chapter 1

The (Mis)use of Information in Decentralised Markets

Abstract

A seller offers an asset in a decentralised market. Buyers have private signals about their common value. I study whether the market becomes allocatively more efficient with (i) more buyers, (ii) better-informed buyers. Both increase the information available about buyers' common value, but also the adverse selection each buyer faces. With more buyers, trade surplus eventually increases and converges to the full-information upper bound if and only if the likelihood ratios of buyers' signals are unbounded from above. Otherwise, it eventually decreases and converges to the no-information lower bound. With better information about trades buyers would have accepted, trade surplus increases. With better information about trades they would have rejected, trade surplus decreases—unless adverse selection is irrelevant. For binary signals, a sharper characterisation emerges: stronger good news increase total surplus, but stronger bad news eventually decrease it.

*This chapter is version 4 of arXiv working paper 2506.06848 (https://arxiv.org/abs/2506.06848).

1.1 Introduction

In this paper, I ask whether more information about buyers' common value for an asset improves or harms allocative efficiency in a decentralised market. The setting is parsimonious: the seller sequentially visits n buyers until one accepts to trade at the seller's commonly known reservation value.[1] Negotiations are private: no buyer knows how many others the seller visited already. Trade is efficient when buyers' common value for the asset—its *quality*—is High, but not when it is Low. Each buyer holds a private signal about the asset's quality; conditional on quality, these signals are IID. A buyer accepts trade when she expects it to yield positive surplus; otherwise she rejects. I ask:

1. Does the expected surplus from trade increase with more buyers, each with an additional signal?
2. Does the expected surplus from trade increase with better-informed buyers, i.e., each with a more informative signal?

Both more and better-informed buyers increase the amount of information available in the market about the asset's quality. However, more information in the market—through either channel—is a double-edged sword for allocative efficiency. On the one hand, it might push buyers to better trades by helping them screen the asset's quality better. On the other hand, it might push them to worse trades by exposing them to greater adverse selection: when there are more buyers in the market, more might have already rejected the seller; when each buyer is better-informed, each rejection might owe to a worse signal. This paper shows that the *kind* of information in the market determines how this trade-off is resolved.

My first main result, Theorem 1.1, answers how increasing the number of buyers in the market influences allocative efficiency. This hinges on whether the likelihood ratios of buyers' signals are unbounded from above. If they are, the expected surplus from trade eventually increases in the number of buyers and converges to the full-information benchmark: a High quality seller almost surely trades, but a Low quality seller never does.

[1] In Section 1.7, I show that this simplifying assumption is without loss in a setting where buyers extend take-it-or-leave-it offers to the seller, and the seller takes an offer unless he expects to secure greater surplus in later visits.

This is the upper bound for equilibrium surplus in the market: all gains from trade are exhausted. If they are not, the expected surplus from trade eventually decreases in the number of buyers and converges to the no-information benchmark: either the seller almost surely trades regardless of his quality, or the expected surplus is zero when he trades. This is the lower bound for equilibrium surplus in the market (Proposition 1.3): buyers' ability to screen the asset's value generates no additional gains from trade.

That the outcome (whether trade occurs) in a large market reveals buyers' common value for the asset if and only if the likelihood ratios of their signals are unbounded from above is reminiscent of a large sealed bid common value auction à la Wilson (1977) and Milgrom (1979). There, too, the outcome (the winning bid) reveals the asset's quality if and only if the bidders' signal structure satisfies the same condition.[2] However, when this condition is violated, trade in a decentralised market either becomes completely uninformative about the asset's quality, or only reveals that the expected gains from trade are zero. In contrast, the winning bid in a large auction may still approximate buyers' common value well.[3]

This also offers an interesting contrast with Lauermann and Wolinsky (2016). In a model like mine[4] but where the gains from trade are always positive, they find that (generically) the outcome in a large market either fully reveals or is completely independent of the asset's quality. Theorem 1.1 establishes that another possibility emerges when, ex-ante, the expected gains from trade are negative: trade might be *partially* informative about the asset's quality. However, trade only reveals the expected gains not to be negative but zero instead[5]—this information has no bearing on total surplus.

In Theorems 1.2 and 1.3, I answer how giving better information to each existing buyer influences allocative efficiency. Theorem 1.2 shows that, when buyers' signal structure is binary, stronger good news (higher likelihood ratio at the top) always increases surplus; but stronger bad news (lower likelihood ratio at the bottom) eventually decreases it. The former might prevent a seller from trading, but recovers surplus in doing so. The latter might

[2]This is when bidders' common value for the item can assume two values. Milgrom (1979) identifies "distinguishability" as a condition that generalises "likelihood ratios unbounded from above" when buyers' common value for the item can assume any number of finite values.

[3]See, for instance, Section IV in Lauermann and Wolinsky (2017).

[4]See the Related Literature section for a more detailed discussion.

[5]Section 1.8 illustrates this with a numerical example.

help a seller trade, but this eventually destroys surplus due to adverse selection. Theorem 1.3 generalises this insight to arbitrary finite signal structures: additional information where a buyer would have accepted trade (a *negative override*) increases surplus; but additional information where she would have rejected trade (a *positive override*) decreases it—unless *adverse selection is irrelevant* in the appropriate sense (Definition 1.5).

To understand the main insight, consider buyers with a binary signal structure: each buyer receives either a good, or a bad signal. For simplicity, ignore equilibrium considerations; simply let buyers accept upon a good signal and reject upon a bad signal. Now, consider revealing additional information to each buyer—another binary signal. This additional information could serve two purposes. If it is revealed after an initial good signal, it could lead the buyer to revise her initial decision to a rejection. I call information that serves this purpose a *negative override*.[6] If it is revealed after an initial bad signal, it could lead her to revise her initial decision to an acceptance. I call such information a *positive override*. In this simple binary-on-binary example, we can interpret a negative override as a strengthening of good news: a buyer can rely on two good signals to accept. A positive override strengthens bad news: a buyer can rely on two bad signals to reject.

A negative override increases surplus. It makes it harder for the seller to trade—a seller some buyer would have accepted before the negative override became available might now be rejected by every buyer. But when this happens, it reveals the expected surplus from trade to be negative: each buyer observed a bad signal and concluded that trading would reduce surplus, despite not knowing (but suspecting) that all buyers reached the same conclusion.

In contrast, a positive override might decrease surplus. It makes it easier for the seller to trade—a seller every buyer would have rejected might now trade with some buyer. However, due to adverse selection, the expected surplus from such a trade might be negative: the buyer who trades with the seller does not observe how many others rejected him previously. If too many did, those buyers' bad signals might reveal expected surplus from trade to be negative despite her good signal. The buyer might find that she traded when she should

[6] I follow the language used in credit markets: a negative (downgrade) override occurs when a human evaluator revises a prospective borrower's algorithmic credit score downwards, in light of overlooked information. A positive (upgrade) override occurs when she revises it upwards. See Section 2.5 in Van Gestel and Baesens (2008); as well as par. 110 in ECB (2024) and pg. 140 in Stellantis Financial Services Italia S.p.A. (2024).

not have.

I show that adverse selection severely limits positive overrides from raising surplus: unless *adverse selection is irrelevant*, i.e., a buyer need not care about the number of previous refusals the seller received, a positive override reduces total surplus.

This insight underpins Theorem 1.2's sharp characterisation for binary signals. To extend it to arbitrary signal structures in Theorem 1.3, I formalise a positive (and, negative) override as a *local mean preserving spread*[7] of a signal upon which buyers reject (and, accept). Studying informativeness at the level of local spreads is essential to the tractability of my exercise but sacrifices no generality: any Blackwell improvement is a combination of finitely many local spreads.

Theorems 1.2 and 1.3 show that too much information can be detrimental for allocative efficiency. So, finally, I study how a regulator can coarsen buyers' information to maximise expected surplus. In Section 1.6, I show that through this policy tool, the regulator aims to prevent a buyer from trading unless adverse selection is irrelevant, i.e., unless she should trade even if everyone else rejected the seller. The implication is striking: the regulator wants buyers to base their decisions on the *highest* number of rejections the seller may have received, not the expected number of rejections.

Contribution

I view the main contribution of my paper to be twofold. First, I study a question that has been largely overlooked by the literature on information aggregation in markets. Most of this literature[8] asks whether the outcome in a large market reflects all information its participants have. Instead, I ask whether a *finite* decentralised market can convert *more information* among its participants to more efficient outcomes. I study two channels which increase information in the market. The first is through an additional buyer, bringing an additional signal to the market. This paper is the first to explore this channel in a decentralised market.[9] The second is through better-informed buyers. To the best of my

[7] See Definitions 1.2 and 1.4.

[8] Prominent and related work in this literature includes Wilson (1977), Milgrom (1979), Riordan (1993) and Lauermann and Wolinsky (2017) for auctions (centralised markets) and Wolinsky (1990), Zhu (2012), and Lauermann and Wolinsky (2016) for decentralised markets.

[9] Riordan (1993) studies how allocative efficiency in a common value auction changes with an additional bidder.

knowledge, this channel has not been explored by previous work.[10]

Second, my findings have important policy implications for markets where trades are negotiated bilaterally and with little to no trade transparency: such as over-the-counter markets,[11] credit markets,[12] and the housing market.[13] Recent technological advances have allowed participants in these markets to enjoy increasingly greater access to information.[14] It is commonly presumed that the "more efficient processing of information, for example in credit markets, financial markets, [...] contribute to a more efficient financial system" (Financial Stability Board 2017). My Theorems 1.2 and 1.3 show that this presumption—which ignores the adverse selection problem in these markets—is misleading: adverse selection may claw back on market participants' ability to screen for efficient trades, and lead to *lower* surplus when each can access better information. My Theorem 1.1 shows that adverse selection might cause increased competition to hurt efficiency, too. This validates a concern empirically recognised by regulators and industry leaders.[15]

Existing regulation in credit markets already limits the information lenders can use to assess borrowers.[16] Directly resonating with Theorems 1.2 and 1.3, ECB guidelines (Feb. 2024) state that "institutions should be more restrictive with positive overrides than with negative ones". I offer a novel justification for such policies, rooted in adverse selection.[17] In Section 1.6, I study how a regulator can design restrictions on market participants'

[10]Notably, Glode and Opp (2019) show that a decentralised OTC market provides buyers with greater information acquisition incentives than a centralised limit-order market, so might be more efficient than the latter. I discuss their work in light of my contribution under Related Literature.

[11]OTC markets are characterised by sequential contacts and little transparency (Duffie 2012; Zhu 2012). Liquidity providers typically make ultimatum offers that last "as long as the breath is warm" (Bessembinder and Maxwell 2008).

[12]In the US and the UK, credit scores mask borrowers' recent applications, and borrowers exercise little bargaining power against lenders (Agarwal et al. 2024 and Consumer Rights 2024).

[13]In the housing market, "buyers and sellers must search for each other" (Han and Strange 2015). Sellers frequently relist, making it difficult to infer how many viewings resulted in no trade: RE/MAX (2024) advises "if a property has been sitting on the market and going stale, there is no harm in relisting it so that it appears fresh and new".

[14]Hedge funds and broker-dealers use increasingly sophisticated data and algorithms to assess trades' profitability (Financial Stability Board 2017); lenders use cutting-edge ML technology in credit scoring (Financial Stability Board 2017); algorithmic traders in housing markets analyse and execute trades faster than traditional investors (Raymond 2024).

[15]Regulators (partially) blamed adverse selection for the collapse of a British bank, HBOS: "the borrowers who came through its doors inevitably included many whom better established banks had turned away" (Kay (2024)).

[16]For instance, following the 2008 crash, the Basel III Accord severely limited the use of "advanced internal ratings systems" to determine credit risk exposure. This overturned the conventions set in Basel II. See BCBS (2017).

[17]Currently, these policies are mostly justified by a distrust in the "robustness and prudence" of lenders' abilities to screen borrowers (BCBS 2017).

information to increase total surplus in the market.

Literature Review

The first question I ask is whether allocative efficiency in a decentralised market increases with more buyers. Riordan (1993) asks this question in a first price auction with common values.[18] There, the adverse selection problem is simply the winner's curse—the winner understands that she had the highest signal among all bidders. In contrast, here, a buyer who trades understands that she had the highest signal among those buyers the seller previously visited. Consequently, the sufficient condition Riordan (1993) identifies[19] for surplus to be increasing or decreasing in the number of bidders differs from the necessary and sufficient condition Theorem 1.1 supplies for a decentralised market.

My second question—whether efficiency increases when buyers are *better-informed*—is novel in this literature. The closest papers, to the best of my knowledge, are Levin (2001) and Glode and Opp (2019). Levin (2001) asks whether in a lemons market à la Akerlof (1970)—where the seller's private information is the root cause of inefficiency—a better-informed seller necessarily hurts efficiency. He finds that the right kind of information can increase efficiency. His environment, and therefore his characterisation, differs from mine. Glode and Opp (2019) show that buyers might have greater incentives to acquire information in an OTC market than in a limit-order market; so, the former can be more efficient than the latter. They investigate a particular information technology: buyers invest in their probability of getting a fully revealing signal. My Theorems 1.2 and 1.3 show that in general, a market with better-informed buyers might be less efficient due to adverse selection.

My model is closest to Zhu (2012) and Lauermann and Wolinsky (2016). Zhu (2012) assumes that trade is efficient regardless of the asset's quality;[20] otherwise his model is identical to mine. He shows that unless the likelihood ratios of buyers' signals are unbounded from above, a large market might fail to be efficient—a High quality seller might

[18] Relatedly, Di Tillio, Ottaviani, and Sørensen (2021) study whether the winning bid in a first price common value auction becomes more informative about the asset's value when there are more bidders. Efficiency is of no direct concern in their setting: trade is always efficient and always materialises.

[19] Where $F_\theta(s)$ is the CDF of the signal distribution for quality $\theta \in \{L, H\}$, he finds that the sign of the expression $\frac{f_H(.)/F_H(.)}{f_L(.)/F_L(.)} - \frac{F_H(.)}{F_L(.)}$ over the support is a sufficient condition to determine this.

[20] Both buyers' common value and the seller's reservation value for a Low quality asset is 0.

fail to trade. Where trading with a Low quality seller is inefficient, Theorem 1.1 offers a stronger conclusion: unless the same condition holds, surplus in a large market converges to the no-information lower bound.

Lauermann and Wolinsky (2016), too, study a decentralised market for a common value asset; but they focus on a large market where the seller (there, the buyer) (i) pays a small cost for each buyer (there, seller) he visits and (ii) has bargaining power. Importantly, there is no efficiency concern: trade is always efficient and executed. They find that generically, the transaction price carries either full or no information about the asset's common value. Furthermore, the seller's costly search further impedes the revelation of the asset's value: the condition necessary for the transaction price to be fully revealing is stronger than the unboundedness of likelihood ratios.

Chen, Ekmekci, and Lauermann (2020) study a similar model to mine. In their model, an organiser sequentially contacts voters, searching for a veto. Each contact is costly. Voters wish the reform to be vetoed in state α, but not in state β. The organiser knows the state, but the voters do not. Chen, Ekmekci, and Lauermann (2020) show that as the number of voters rises, voters' payoffs converge to the full-information benchmark: the reform is nearly always vetoed in state α, but passed in state β. This is achieved through the organiser's incentives to continue sampling in state α, but to stop sampling in state β.

Besides having an organiser who samples at cost rather than a seller who samples freely, Chen, Ekmekci, and Lauermann (2020) differs crucially from this paper in their voters' behaviour. Their voters have common interests: they care about whether *some* voter vetoes the policy. So, their behaviour maximises their aggregate payoffs: each voter conditions on being pivotal—the last voter who may refuse a veto. In my model, a buyer cares only about whether *she* buys the asset. Her behaviour maximises her individual payoff: she conditions on the expected number of buyers who refused the seller before her. This acts as an additional friction which prevents the market from achieving higher efficiency.

Cavounidis (2022) also studies a model very similar to mine; but there, "appraisers" (replacing the buyers here) must acquire costly information about the state. Rather than focusing on the effect of increasing the amount of information available in the market, Cavounidis (2022) asks how many appraisers the market can profitably accommodate, given

those appraisers' information acquisition costs. He finds that when the cost of information is "divisible", the market can accommodate arbitrarily many appraisers; but when it is "chunky", the size of the market is constrained by the adverse selection problem this paper also highlights.

My model also admits a social learning interpretation, in the tradition of Bikhchandani, Hirshleifer, and Welch (1992). It can be considered as a variant of the classic model: later decision makers (buyers) are called to decide only if those before them reject, and no one observes her position in the queue (as in Herrera and Hörner (2013)). Most work in this literature focuses on whether full learning attains with a large number of decision makers. Instead, my results speak to how more information, through more or better-informed decision makers, influences the welfare of finitely many decision makers.

The remainder is organised as follows. Section 1.2 presents the model. Section 1.3 presents preliminary analyses about equilibria and total surplus in the market. Section 1.4 presents Theorem 1.1. Section 1.5 presents Theorems 1.2 and 1.3 (in Subsections 1.5.1 and 1.5.2). Section 1.6 discusses how buyers' information can be coarsened to maximise total surplus. Section 1.7 presents an extension where buyers offer take-it-or-leave it prices to the seller. Section 1.8 presents a numerical example that supplements the discussion in Section 1.4. Section 1.9 presents the proofs and results the main text omits.

1.2 Model

The seller (he) of an indivisible asset sequentially visits $n \in \mathbb{N}$ prospective buyers (she) in a uniformly random order. He sells to the first one who accepts to pay his reservation value $c \in [0, 1]$. The asset's (seller's) quality θ is either High or Low, $\theta \in \{H, L\}$. When the buyer he visits rejects trade, no transaction takes place and the seller proceeds to visit the next buyer. If the buyer accepts, she pays the seller his reservation value and receives the asset. She enjoys a return of 1 if the asset's quality is High, but 0 if it is Low. The game ends when a buyer accepts the seller, or they all reject him.

At the outset of the game, the asset's quality is unknown;[21] all players share the common prior that it is High with probability ρ and Low otherwise. Each buyer obtains

[21] Until Section 1.7, the seller's knowledge about the asset's quality is immaterial.

additional private information about the asset's quality through the outcome of a Blackwell experiment $\mathcal{E} = (\mathbf{S}, p_L, p_H)$. The outcome s of the experiment—the buyer's *signal*—is drawn from the finite set $\mathbf{S}$ with a distribution p_θ. Conditional on the asset's quality, buyers' signals are IID. The joint distribution of buyers' signals conditional on the asset's quality is common knowledge.

The buyer visited by the seller receives no information about how many others the seller previously visited. Nonetheless, she deduces that all those buyers rejected the seller. Through this, she extracts additional information about the asset's quality.

The buyer forms her posterior belief about the asset's quality using the information conveyed by the seller's visit and her private signal. First, she uses the information conveyed by the seller's visit to revise her prior belief ρ to an interim belief ψ. Then, she uses her private signal to revise her interim belief to a posterior belief $\mathbb{P}_\psi(\theta = H \mid s)$.

A buyer's strategy $\sigma : \mathbf{S} \to [0, 1]$ maps every signal $s \in \mathbf{S}$ she might observe to a probability $\sigma(s)$ with which she accepts to trade. Her strategy σ is *optimal against the interim belief* ψ if, given this interim belief and her signal, the buyer accepts (rejects) to trade whenever her expected payoff from trading with the seller is positive (negative):

$$\sigma(s) = \begin{cases} 0 & \mathbb{P}_\psi(\theta = H \mid s) < c \\ 1 & \mathbb{P}_\psi(\theta = H \mid s) > c \end{cases}$$

She may accept to trade with any probability when she expects zero surplus from trading.

I focus on *symmetric Bayesian Nash Equilibria* of this game. Hereafter, I reserve the term *equilibrium* for such equilibria unless I state otherwise. An *equilibrium* is a strategy and interim belief pair (σ^*, ψ^*) such that:

1. The interim belief ψ^* is *consistent* with the strategy σ^*; i.e., it is the interim belief of a buyer who believes all other buyers use the strategy σ^*.
2. The strategy σ^* is optimal given the interim belief ψ^*.

I call any strategy σ^* that constitutes part of an equilibrium an *equilibrium strategy.*

1.3 Buyers' Beliefs, Equilibria, and Total Surplus

This section lays the necessary groundwork to discuss my main results. First, I discuss how buyers form their interim beliefs, and the fundamental properties of the set of equilibria. Then, I discuss how the total surplus from trade varies across different equilibria.

1.3.1 Buyers' Beliefs and Equilibria

No buyer learns how many others the seller visited before her. But, she deduces that all those past visits resulted in rejections. How does she interpret this information?

When each buyer uses a strategy σ, a seller of quality θ faces a probability $r_\theta(\sigma;\mathcal{E})$ of getting rejected in any of his visits:

$$r_\theta(\sigma;\mathcal{E}) := 1 - \sum_{j=1}^{m} p_\theta\left(s_j\right) \times \sigma\left(s_j\right)$$

Every buyer understands that the seller is equally likely to decide to visit any number $k \in \{0,1,2,...,n-1\}$ of other buyers before her. She will receive the seller's visit if and only if he is rejected by all those k buyers. Therefore, she assigns a probability $\nu_\theta\left(\sigma;\mathcal{E}\right)$ to being visited by the seller:

$$\nu_\theta\left(\sigma;\mathcal{E}\right) := \frac{1}{n} \times \sum_{k=0}^{n-1} r_\theta\left(\sigma;\mathcal{E}\right)^k$$

When the seller *does* visit her, the buyer uses this information to update her prior belief about the seller's quality to an interim belief ψ:

$$\begin{aligned}\psi = \mathbb{P}\left(\theta = H \mid \text{visit received}\right) &= \frac{\mathbb{P}\left(\text{visit received} \mid \theta = H\right) \times \mathbb{P}(\theta = H)}{\mathbb{P}\left(\text{visit received}\right)} \\ &= \frac{\rho \times \nu_H\left(\sigma;\mathcal{E}\right)}{\rho \times \nu_H\left(\sigma;\mathcal{E}\right) + (1-\rho) \times \nu_L\left(\sigma;\mathcal{E}\right)}\end{aligned}$$

This is the unique interim belief that is *consistent* with every buyer using the strategy σ. The buyer then uses her private signal $s \in \mathbf{S}$ about the seller's quality to update her

interim belief to a posterior belief:

$$\mathbb{P}_{\psi}\left(\theta = H \mid s\right) = \frac{\psi \times p_H(s)}{\psi \times p_H(s) + (1-\psi) \times p_L(s)}$$

Note that the informational content of the buyer's signal $s \in \mathbf{S}$ is distilled by the ratio $\frac{p_H(s)}{p_H(s)+p_L(s)}$. For notational convenience, I will use the signal's label, s, to refer to this ratio:

$$s := \frac{p_H(s)}{p_H(s) + p_L(s)} \in [0,1] \qquad \text{for all } s \in \mathbf{S}$$

Under this notation, the ratio $\frac{s}{1-s}$ simply equals the signal's likelihood ratio, $\frac{p_H(s)}{p_L(s)}$. For further convenience, I also enumerate the signals $\mathbf{S}$ in order of increasing likelihood ratios; $\mathbf{S} := \{s_1, s_2, ..., s_m\}$ where $s_1 \leq s_2 \leq ... \leq s_m$. Note that, for the same interim belief, a buyer's posterior belief is increasing in her signal's index.

Whenever a strategy σ^* is optimal against the unique interim belief ψ^* consistent with it, the pair (σ^*, ψ^*) forms an equilibrium. In principle, there might be many such pairs, or none at all. Proposition 1.1 sets the ground by ruling the latter possibility out and characterising the set of equilibria.

Proposition 1.1. Let Σ be the set of equilibrium strategies. Then:

1. Σ is non-empty and compact.

2. Any equilibrium strategy σ^* is monotone: for any $\sigma^* \in \Sigma$, $\sigma^*(s) > 0$ for some $s \in \mathbf{S}$ implies that $\sigma^*(s') = 1$ for every $s' \in \mathbf{S}'$ such that $s' > s$.

3. All equilibria exhibit adverse selection: $\psi^* \leq \rho$ for any interim belief ψ^* consistent with an equilibrium strategy $\sigma^* \in \Sigma$.

Proof outline: To establish the existence of an equilibrium, I construct a best response correspondence: Φ for buyers. Φ maps any strategy σ to the set of strategies that are optimal against the unique interim belief consistent with σ; i.e. those that maximise a buyer's expected payoff when her all her peers use the strategy σ. Note that σ^* is an equilibrium strategy if and only if it is a fixed point of this best response correspondence; i.e., $\sigma^* \in \Phi(\sigma^*)$. Through a routine application of Kakutani's Fixed Point Theorem, I show

that Φ indeed has a fixed point. In the process, I prove that Φ is upper semicontinuous; this also establishes that the set of equilibrium strategies is compact.

Monotonicity is a straightforward necessity for a strategy to be optimal: higher signals induce higher posterior beliefs, so buyers (weakly) prefer to trade upon higher signals. A crucial consequence of monotonicity is that a Low quality seller is likelier to be rejected in any of his visits, as buyers are likelier to observe lower signals for him. Thus, the seller is adversely selected through his past rejections, and buyers' interim beliefs always lie below their prior beliefs. □

An equilibrium is guaranteed to exist, but it need not be unique. The following example, which I will modify and revisit on occasion, illustrates this. Let there be two buyers who share the prior belief $\rho = 0.5$, and the seller's reservation value be $c = 0.2$. Furthermore, let buyers' experiment $\mathcal{E}$ be binary, $\mathbf{S} = \{0.2, 0.8\}$, and its outcome have the conditional distribution:

$$p_L(s) = \begin{cases} 0.8 & s = 0.2 \\ 0.2 & s = 0.8 \end{cases} \qquad p_H(s) = \begin{cases} 0.2 & s = 0.2 \\ 0.8 & s = 0.8 \end{cases}$$

There are two equilibrium strategies in this example,[22] which I denote as $\hat{\sigma}$ and $\check{\sigma}$. Table 1.1 summarises these strategies, and how buyers form their interim and posterior beliefs under them.

	$\check{\sigma}$	$\hat{\sigma}$
Prob. the buyer accepts upon s = 0.8	$\check{\sigma}(0.8) = 1$	$\hat{\sigma}(0.8) = 1$
Prob. the buyer accepts upon s = 0.2	$\check{\sigma}(0.2) = 1$	$\hat{\sigma}(0.2) = 0$
Buyers' interim belief	$\psi = 0.5$	$\psi = 0.4$ [23]
Buyers' posterior belief upon s = 0.8	0.8	0.7
Buyers' posterior belief upon s = 0.2	0.2	≈ 0.14

Table 1.1: Running Example: Comparing Equilibrium Strategies $\check{\sigma}$ and $\hat{\sigma}$

[22] There are no other equilibria. Under any monotone strategy that assigns a positive probability to trade, buyers' interim belief lies between 0.5 and 0.4—so buyers must always trade upon the high signal 0.8. Lemma 1.17 in Section 1.9.3 shows that in equilibrium, buyers either always or never trade upon the low signal.

[23] The interim belief in this case is easily calculated as: $\psi = \frac{1+r_H(\sigma;\mathcal{E})}{(1+r_H(\sigma;\mathcal{E}))+(1+r_L(\sigma;\mathcal{E}))} = \frac{1.2}{1.2+1.8} = 0.4$.

Under the strategy $\check{\sigma}$, a buyer accepts trade regardless of her signal. This eliminates adverse selection: since the first buyer the seller visits accepts trade, whoever he visits is certain that she is the first one he visited; and so, buyers' interim belief ψ equals their prior ρ. In this equilibrium, a buyer finds trade optimal even if she receives the low signal 0.2; given the posterior belief this signal induces, she expects zero net surplus from trade.

Under the strategy $\hat{\sigma}$, a buyer accepts trade if only if she receives the high signal 0.8. Buyers' selectivity triggers adverse selection: each buyer understands that she need not be the first one the seller visited. So, buyers' interim belief ψ falls below their prior belief. A buyer no longer finds trade optimal when she receives the low signal 0.2: this signal induces a posterior belief ≈ 0.14, so she expects a loss. She does, however, find it optimal when she receives high signal $s = 0.8$: this signal induces a posterior belief of 0.7, so she expects positive net surplus from trade.

In the equilibrium where buyers use the strategy $\hat{\sigma}$, the seller is likelier to be rejected in any of his visits. The buyers are *more selective*—they are (weakly) likelier to reject trade at any signal they might observe.

Definition 1.1. Where σ' and σ are two strategies, σ' is *more selective than* σ (or, σ is *less selective than* σ') if $\sigma'(s) \leq \sigma(s)$ for all $s \in \mathbf{S}$.

Selectivity offers a natural way to order buyers' equilibrium strategies. Proposition 1.2 shows that it is also a complete order over them.

Proposition 1.2. Selectivity is a complete order over the set of equilibrium strategies Σ. Moreover, Σ contains a *most* and *least* selective strategy, $\hat{\sigma} \in \Sigma$ and $\check{\sigma} \in \Sigma$ respectively:

$$\hat{\sigma}(s) \leq \sigma^*(s) \leq \check{\sigma}(s) \qquad \text{for all } s \in \mathbf{S} \text{ and } \sigma^* \in \Sigma$$

Proof. By Proposition 1.1, the set of equilibrium strategies Σ is a subset of the set of monotone strategies. The latter is a chain under the selectivity order: for any signal $s \in \mathbf{S}$ and two monotone strategies σ and σ', we have:

$$\sigma'(s) > \sigma(s) \implies \begin{array}{ll} 1 = \sigma'(s^{\cdot}) \geq \sigma(s^{\cdot}) & \text{for any } s^{\cdot} > s \in \mathbf{S} \\ \sigma'(s_{\cdot}) \geq \sigma(s_{\cdot}) = 0 & \text{for any } s_{\cdot} < s \in \mathbf{S} \end{array}$$

Since any subset of a chain is also a chain, Σ is a chain too.

By Proposition 1.1, Σ is a compact set. Since it is also a chain, by applying a suitably general Extreme Value Theorem (see Theorem 27.4 in Munkres (2000)) to the identity mapping on Σ, one verifies that Σ has a minimum and maximum element with respect to this order; i.e. there are two strategies $\hat{\sigma}, \check{\sigma} \in \Sigma$ such that for any other strategy $\sigma^* \in \Sigma$ we have $\hat{\sigma}(s) \leq \sigma^*(s) \leq \check{\sigma}(s)$ for all $s \in \mathbf{S}$. □

I call an equilibrium "more selective than another" whenever buyers use a more selective strategy in the former.

1.3.2 Total Surplus

Trading generates a surplus of $1-c$ when the asset's quality is High, but destroys a surplus of c when the asset's quality is Low. So, the expected surplus from trade in the market—total surplus, for short—depends on how well buyers can screen the asset's quality before they decide whether to trade. Given buyers' experiment $\mathcal{E}$ and strategy σ, total surplus equals:

$$\begin{aligned}\Pi(\sigma;\mathcal{E}) &:= (1-c) \times \mathbb{P}_{\sigma;\mathcal{E}}\left(\theta = H \cap \text{some buyer trades}\right) - c \times \mathbb{P}_{\sigma;\mathcal{E}}\left(\theta = L \cap \text{some buyer trades}\right) \\ &= (1-c) \times \rho \times [1 - r_H(\sigma;\mathcal{E})^n] - c \times (1-\rho) \times [1 - r_L(\sigma;\mathcal{E})^n]\end{aligned}$$

Buyers fully appropriate this surplus—the seller is only paid his reservation value when he trades.

Two benchmarks are natural to consider. If buyers had full information about the asset's quality, they would trade whenever the asset's quality is High, but never when it is Low; all gains from trade would be realised. Total surplus in this *full-information benchmark* would equal $\Pi^f := \rho \times (1-c)$.

If instead, buyers had no information about the asset's quality, their decisions would be guided solely by their prior belief. There would be no adverse selection: previous rejections would convey no private information since no buyer has any. A buyer would trade if, given her prior belief, she expected positive surplus from doing so, and reject otherwise. Total surplus in this *no-information benchmark* would equal $\Pi^n := \max\{0, \rho - c\}$.

Proposition 1.3 establishes that, in equilibrium, total surplus is bounded by these benchmarks; moreover, total surplus is always higher in more selective equilibria.

Proposition 1.3. Equilibrium total surplus is bounded above by the full-information benchmark Π^f and below by the no-information benchmark $\Pi^\emptyset$. Furthermore, it is higher under more selective equilibrium strategies:

$$\max\{0, \rho - c\} = \Pi^\emptyset \leq \Pi(\sigma^{**}; \mathcal{E}) \leq \Pi(\sigma^*; \mathcal{E}) \leq \Pi^f = \rho \times (1 - c)$$

where σ^* and σ^{**} are two equilibrium strategies such that σ^{**} is more selective than σ^*.

Proof. The second part of the Proposition follows as a corollary to Lemma 1.11 in Section 1.9.2, where I prove that total surplus decreases whenever buyers jointly deviate from an equilibrium strategy to a less selective monotone strategy.

For the first part of Proposition 1.3, note that all gains from trade is realised when buyers have full information; hence Π^f bounds total surplus from above. Since total surplus equals buyers' surplus from trade, it is bounded below by 0 in any equilibrium—a buyer can always reject. Thus, when $\rho \leq c$, $\Pi^\emptyset = 0$ bounds total surplus from below. Now let $\rho > c$, and assume for contradiction that there is an equilibrium strategy σ^* such that $\Pi(\sigma^*; \mathcal{E}) < \Pi^\emptyset$. Then:

$$\begin{aligned}
&\mathbb{P}_{\sigma^*;\mathcal{E}}(\text{some buyer trades}) \times [\mathbb{P}_{\sigma^*;\mathcal{E}}(\theta = H \mid \text{some buyer trades}) - c] < \\
&\mathbb{P}_{\sigma^*;\mathcal{E}}(\text{some buyer trades}) \times [\mathbb{P}_{\sigma^*;\mathcal{E}}(\theta = H \mid \text{some buyer trades}) - c] + \\
&\qquad \mathbb{P}_{\sigma^*;\mathcal{E}}(\text{no buyer trades}) \times [\mathbb{P}_{\sigma^*;\mathcal{E}}(\theta = H \mid \text{no buyer trades}) - c] = \Pi^\emptyset
\end{aligned}$$

So, $\mathbb{P}_{\sigma^*;\mathcal{E}}(\theta = H \mid \text{no buyer trades}) > c$. However, σ^* then cannot be an equilibrium strategy; each buyer has a profitable deviation to trade with the seller whenever he visits. □

Buyers' private information helps them raise surplus by avoiding trade when the asset's quality is Low, and executing it when it is High. Equilibrium multiplicity presents a trade-off: in a more selective equilibrium, trade is less likely—this conserves surplus when the asset's quality is Low. In a less selective equilibrium, trade is more likely—this raises surplus when the asset's quality is High. Proposition 1.3 establishes that this trade-off is always resolved in favour of more selective equilibria.

The intuition is illustrated by our running example. There, we identified two equilibrium strategies, $\hat{\sigma}$ and $\check{\sigma}$. These are the only equilibrium strategies; so, $\hat{\sigma}$ is the most

selective equilibrium strategy in that example, and $\check{\sigma}$ is the least selective. Table 1.2 summarises how the probability that the seller trades and the total surplus vary across these equilibria, and compares them to the full- and no-information benchmarks.

	no-info.	$\check{\sigma}$	$\hat{\sigma}$	**full-info.**
Prob. seller trades when $\theta = \mathbf{H}$	1	1	0.96 [24]	1
Prob. seller trades when $\theta = \mathbf{L}$	1	1	0.36	0
Total surplus	0.3	0.3	0.348	0.4

Table 1.2: Running Example: Comparing Surplus Across Equilibria

Unless both buyers observe the low signal 0.2, the seller trades under both equilibria. If both observe the low signal, both buyers reject the seller in the most selective equilibrium: each fears that the other buyer might have already rejected him, which would indicate that that buyer also observed a low signal. The fact that neither buyer wants to trade in the event that both observe a low signal reveals that the expected surplus from trading with such a seller is negative. Nonetheless, this seller trades in the least selective equilibrium: now, each buyer understands that the seller visits her only if he has not visited the other yet, so she becomes more optimistic about the other buyer's signal. This results in lower total surplus.

1.4 Efficiency with More Buyers

In this section, I discuss how total surplus changes as the number of buyers in the market increases.

Theorem 1.1. Let $\Pi^n(\hat{\sigma}; \mathcal{E})$ be total surplus under the most selective equilibrium in a market with n buyers. If $\mathcal{E}$ has an outcome that fully reveals High quality ($s_m = 1$), the sequence $\{\Pi^n(\hat{\sigma}; \mathcal{E})\}_{n=1}^{\infty}$ is eventually increasing and converges to surplus in the full-information benchmark. Otherwise, it is eventually decreasing and converges to surplus in the no-information benchmark.

Each additional buyer brings an additional signal about the asset's quality to the market. As the market becomes arbitrarily large, buyers' collective information becomes

[24] This is the probability that at least one buyer will receive the high signal 0.8 when the seller has High quality: $1 - p_H^2(0.2) = 0.96$. For a seller of Low quality, this probability is: $1 - p_L^2(0.2) = 0.36$.

sufficient to fully reveal the asset's quality—unless buyers' signals carry no information. Whether the market outcome incorporates this information and reaches full allocative efficiency depends on the *kind* of information each buyer has.

If buyers have a signal which fully reveals High quality, total surplus reaches its full-information upper bound as the number of buyers grows arbitrarily large. Beyond a threshold number of buyers, adverse selection forces each buyer to reject the seller unless she observes that fully-revealing high signal. Therefore, a Low quality seller never trades. A High quality seller, however, might. Moreover, a High quality seller is likelier to trade in a larger market, as it is likelier that at least one buyer will observe the fully-revealing high signal.

If, however, buyers do not have such a signal, the market experiences a surplus breakdown as the number of buyers grows large—total surplus reaches its no-information lower bound. How market outcomes evolve as the number of buyers grows depends on whether the expected gains from trade are positive, $\rho > c$, or weakly negative, $\rho \leq c$:

- When the expected gains from trade are positive, no matter how large the market, a buyer who observes the highest possible signal trades when the seller visits. The larger the market, the likelier that some buyer will observe this signal—regardless of the seller's quality. Beyond a threshold number of buyers, this hurts total surplus through the increased probability that a Low quality seller trades. As the number of buyers grows arbitrarily large, the seller almost surely trades. This yields the level of surplus in the no-information benchmark, $\rho - c$.

- When the expected gains from trade are negative, a buyer accepts the seller with an arbitrarily small probability in a large market, if at all. Adverse selection, however, ensures that she expects zero surplus from doing so. The seller may trade with positive probability, but total surplus is equal to that in the no-information benchmark: 0.

Notably, even when buyers have a signal that fully reveals High quality, total surplus need not converge to the full-information benchmark in *every* equilibrium. To see this, consider a slightly modified version of our running example. As before, the common prior is $\rho = 0.5$, and the seller's reservation value is $c = 0.2$. But we modify buyers' experiment;

they now observe the outcome of $\mathcal{E}^g = \left(\mathbf{S}^g, p_L^g, p_H^g\right)$:

$$p_L^g(s) = \begin{cases} 1 & s = 0.2 \\ 0 & s = 1 \end{cases} \qquad p_H^g(s) = \begin{cases} 0.25 & s = 0.2 \\ 0.75 & s = 1 \end{cases}$$

The signal $s = 1$ fully reveals High quality. However, buyers trade regardless of the signal they receive in the least selective equilibrium—irrespective of the number of buyers in the market. Thus, the seller always trades. Total surplus always equals that in the no-information benchmark, $\Pi^{\emptyset} = 0.3$. On the other hand, buyers only trade upon the high signal $s = 1$ in the most selective equilibrium. Hence, a Low quality seller never trades. A High quality seller trades with probability $1 - (0.25)^n$ in a market with n buyers. In an arbitrarily large market, he trades almost surely; total surplus converges to the full-information benchmark, $\Pi^f = 0.5 \times [1 - 0.2]$.

Theorem 1.1 helps compare a decentralised market with a first price auction where the seller simultaneously solicits buyers' bids and sells to the highest bidder. There too, total surplus may decrease with an additional bidder (Riordan (1993)), due to an increased probability of trade when it is inefficient. Similar to a decentralised market, adverse selection is the culprit: the winner's curse is more severe in an auction with more bidders. However, adverse selection in a decentralised market differs from the winner's curse in an auction: a buyer who trades in a decentralised market understands that her signal was the highest among the *previous* buyers to be visited; a bidder who wins in an auction understands that her signal was the highest among *all* bidders. Thus, Riordan (1993)'s sufficient condition for an additional participant to decrease total surplus differs from the necessary and sufficient condition Theorem 1.1 recovers for total surplus to eventually increase with the number of buyers in a decentralised market.

This condition is also necessary and sufficient for buyers' common value to be revealed through trade in a decentralised market. Strikingly, the same condition is also necessary and sufficient for a first price common value auction to reveal bidders' common value through the winning bid.[25] However, when this condition is violated, a common value

[25]See Wilson (1977) and Milgrom (1979) for this classic result: in a setting where the item's common value may take countably many values, Milgrom (1979) recovers "distinguishability" as a necessary and sufficient condition. When the item's value is binary, "distinguishability" corresponds to an unbounded likelihood ratio from above.

auction with an arbitrarily large number of bidders may nonetheless aggregate information "well".[26] This is not the case in a decentralised market with an arbitrarily large number of buyers. Instead, three possibilities emerge:

1. The expected gains from trade are positive, $\rho > c$, and the seller almost surely trades. The fact that she trades is completely uninformative about the asset's quality.

2. The expected gains from trade are negative, $\rho < c$, and the seller never trades.[27] The fact that she does not trade is completely uninformative about the asset's quality.

3. The expected gains from trade are negative, $\rho < c$, and if a buyer trades, she expects zero surplus from doing so. Trade—when it happens—reveals that the asset has High quality with probability c, i.e., that the expected surplus from trade is zero.[28]

So, whenever the likelihood ratio of buyers' signals is *not* unbounded at the top, trade is at most partially informative about the asset's quality, unlike in an auction. The information it reveals has no bearing on market participants' surplus: at most, trade is revealed to be to be no worse than no trade in expectation.

This offers an interesting contrast with Lauermann and Wolinsky 2016, too.[29] They assume that trade is always efficient, and find that (generically) the outcome in a large market either fully reveals or is completely uninformative about buyers' common value for the asset. Theorem 1.1 echoes their finding where the expected gains from trade are positive, $\rho \geq c$. However, it also shows that another possibility arises when the expected gains from trade are negative, $\rho < c$: trade in a large market may be partially informative about the asset's quality.

1.5 Efficiency with Better-Informed Buyers

In this section, I discuss how giving each buyer better information—a Blackwell more informative experiment—affects total surplus. Throughout, I take the number of buyers n to be a primitive rather than a parameter. Equilibrium surplus attains its extremes in the

[26] See, for instance, Section IV in Lauermann and Wolinsky (2017).
[27] For instance, if buyers' experiment is uninformative about the asset's value.
[28] Section 1.8 illustrates this scenario with a numerical example.
[29] Lauermann and Wolinsky (2017) find a similar result in a first price auction where the number of bidders varies with the common value of the auctioned item.

most and least selective equilibria. My main results describe the comparative statics of total surplus under both of these equilibria. For brevity, I write *equilibrium** wherever the reader may read *the most* or *least selective equilibrium*. I denote the equilibrium* strategy under an experiment $\mathcal{E}$ as $\sigma^\star_{\mathcal{E}}$.

If there were a single buyer in the market—exposed to no adverse selection—a Blackwell improvement of her experiment would be necessary and sufficient for total surplus to rise regardless of the seller's reservation value.[30] The reason is simple: better information improves her ability to screen the asset's quality and target efficient trades.

In a market with multiple buyers, buyers' ability to screen the seller is shaped both by the quality of their private signals and the extent of adverse selection each face in the market. So, better information becomes a double-edged sword. On the one hand, it allows each buyer to screen the seller more effectively. On the other hand, it might exacerbate adverse selection: previous rejections might become likelier, and each might carry worse news about the quality of the asset. This latter channel pushes the buyer to worse trades. If it overwhelms buyers' increased ability to screen the asset's quality, it might lead to lower total surplus. How this trade-off is resolved is determined by the *kind* of improvement in buyers' information.

1.5.1 Binary Experiments

I start by restricting buyers to binary signals. In this setting, I obtain a sharper characterisation and uncover the main insights which drive my result for general signal structures, in Theorem 1.3.

A binary experiment $\mathcal{E}$ has two possible outcomes $s_1, s_2 \in \mathbf{S}$, which I relabel as $s_L, s_H \in \mathbf{S}$ for convenience.[31] The low outcome $s_L \in [0, 0.5]$ decreases a buyer's interim belief about the quality of the asset, while a high outcome $s_H \in [0.5, 1]$ increases it.

Ranking two binary experiments $\mathcal{E}'$ and $\mathcal{E}$ in their (Blackwell) informativeness is a simple exercise. Where the former has the possible outcomes $s'_L, s'_H \in \mathbf{S}'$, the experiment

[30] In general, a Blackwell improvement is sufficient, but not necessary for a decision maker to extract higher value from a decision problem (see Blackwell (1953)). However, it is necessary for the decision maker to extract higher value from *any* decision problem where the unknown state is binary—such our buyers' screening problem. I present a self contained proof of this fact in Section 1.9.2, Lemma 1.12 for completeness.

[31] Recall from Section 1.3 that, to ease notation, I use labelling convention $s := \frac{p_H(s)}{p_H(s)+p_L(s)}$.

$\mathcal{E}'$ is Blackwell more informative than $\mathcal{E}$ if and only if:[32]

- it has a lower likelihood ratio at the bottom, $s'_L \leq s_L$; i.e., delivers *stronger bad news*, and
- it has a greater likelihood ratio at the top, $s'_H \geq s_H$; i.e. delivers *stronger good news*.

I illustrate these two improvements in information in Figures 1.5.2a and 1.5.2c.

Theorem 1.2 (illustrated in Figure 1.5.1) answers how total surplus evolves when buyers' binary experiment becomes Blackwell more informative.

Theorem 1.2. Let buyers' experiment $\mathcal{E}$ be binary. Then, equilibrium* total surplus is increasing in the strength of good news (s_H) but is quasiconcave and eventually decreasing in the strength of bad news (s_L).

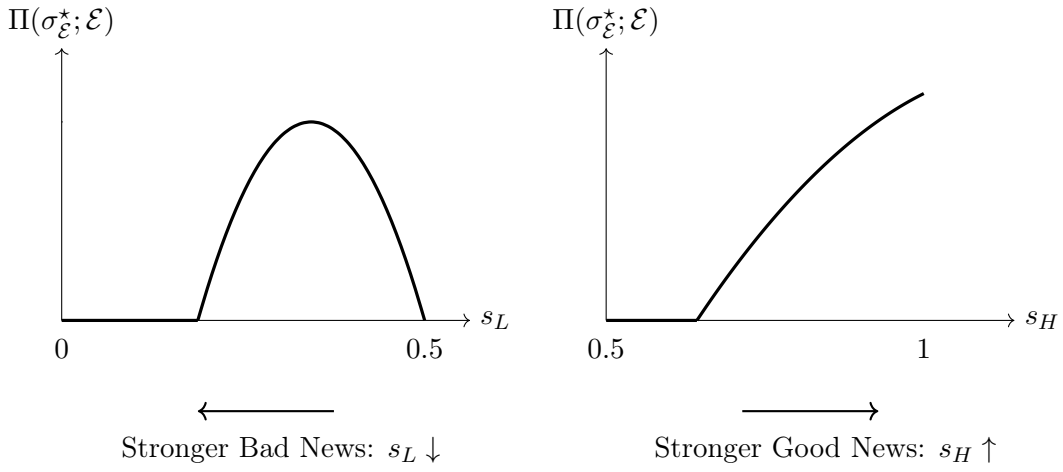

Figure 1.5.1: Theorem 1.2 illustrated

To understand the intuition behind Theorem 1.2, let us start from the case of stronger good news. Instead of the experiment $\mathcal{E}$, buyers now observe the outcome of $\mathcal{E}' = (\mathbf{S}', p'_L, p'_H)$, which delivers stronger good news than $\mathcal{E}$, $s'_H > s_H$, but the same strength of bad news as $\mathcal{E}$, $s_L = s'_L$. Brushing equilibrium considerations aside, simply assume that, under both experiments, a buyer accepts upon the high signal, s_H or s'_H, and rejects upon the low signal, s_L or s'_L. How does, then, this improvement in buyers' information affect the seller's chances of trading?

The answer is clearest when we reinterpret the additional information a buyer can extract from $\mathcal{E}'$ as an additional signal she might observe. Instead of replacing the original

[32] See Section 12.5 in Blackwell and Girshick (1954) for a textbook exposition of this classic result.

experiment $\mathcal{E}$ with $\mathcal{E}'$, imagine that a buyer who observes an initial high signal s_H from the experiment $\mathcal{E}$ then observes the outcome of an *additional* binary experiment $\mathcal{E}^a = (\{s_L^a, s_H^a\}, p_L^a, p_H^a)$. Conditional on the asset's quality, the outcomes of these additional signals are IID across buyers, and independent of the first signal they observe. Their conditional distributions are such that when appended to $\mathcal{E}$ as such, the experiment $\mathcal{E}^a$ mimics the improvement in information that $\mathcal{E}'$ offers:

$$\frac{s_H}{1-s_H} \times \frac{p_H^a\left(s_H^a\right)}{p_L^a\left(s_H^a\right)} = \frac{s'_H}{1-s'_H} \qquad \frac{s_H}{1-s_H} \times \frac{p_H^a\left(s_L^a\right)}{p_L^a\left(s_L^a\right)} = \frac{s'_L}{1-s'_L}$$

A buyer who observes the initial low signal s_L receives no further information. I illustrate this construction in Figure 1.5.2b.

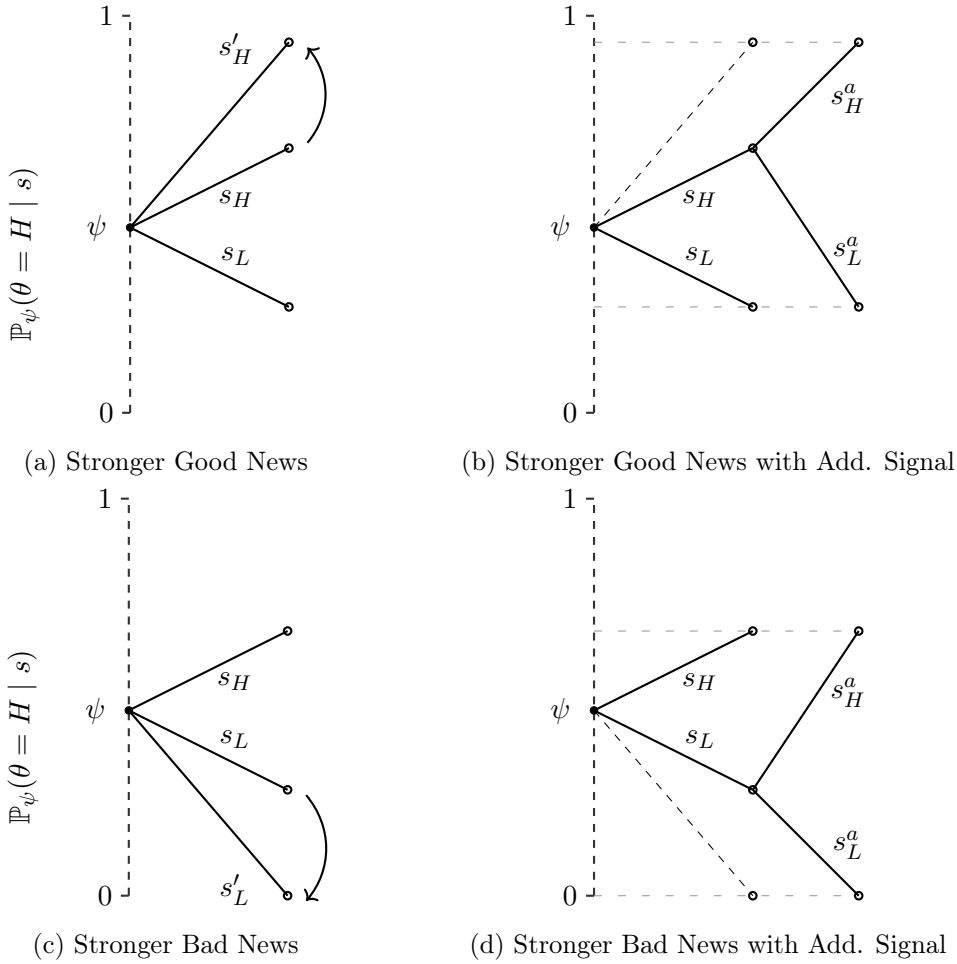

(a) Stronger Good News

(b) Stronger Good News with Add. Signal

(c) Stronger Bad News

(d) Stronger Bad News with Add. Signal

Figure 1.5.2: Blackwell Improvements of a Binary Signal

So, observing the sequence (s_H, s_H^a) conveys the same information as the signal s'_H from $\mathcal{E}'$ does. This information leads to an acceptance. Either the sequence (s_H, s_L^a) or

the signal s_L convey the same information as the signal s'_L from $\mathcal{E}'$ does. This information leads to a rejection.

Our reinterpretation reveals how this additional information affects trade. A buyer who observes an initial low signal s_L receives no additional information. She rejects the seller, as before. However, a buyer who receives an initial high signal s_H receives additional information through $\mathcal{E}^a$. Absent this additional information, she would have traded. But a low signal s^a_L from $\mathcal{E}^a$ "negatively overrides" that initial verdict: now, she rejects the seller.

Thus, stronger good news jeopardises trade: a seller who previously would have traded with some buyer may now be rejected by every buyer. This raises total surplus: the seller is rejected by every buyer because each observe a low signal, either s_L or s^a_L. Trading with such a seller would yield negative surplus in expectation: if a buyer with a low signal expects a negative surplus from trade when she merely *suspects* others to have received low signals, she would certainly expect a negative surplus if she *knew* all others had received low signals too.

Now, let us turn to the case of stronger bad news. To shed light on the threshold beyond which stronger bad news hurts total surplus, let us focus on a an experiment $\mathcal{E}'$ that delivers "marginally" stronger bad news than $\mathcal{E}$, i.e., $s'_L = s_L - \delta$ for a vanishingly small $\delta > 0$, but the same strength of good news as $\mathcal{E}$, $s'_H = s_H$.

As before, let us reinterpret the additional information a buyer can extract from $\mathcal{E}'$ as an additional signal she can observe. This time, she observes this additional signal only after an initial low signal s_L. As before, the conditional distribution of this additional signal is such that when appended to $\mathcal{E}$, it mimics the improvement $\mathcal{E}'$ offers over $\mathcal{E}$:

$$\frac{s_L}{1-s_L} \times \frac{p^a_H\left(s^a_H\right)}{p^a_L\left(s^a_H\right)} = \frac{s_H}{1-s_H} \qquad \frac{s_L}{1-s_L} \times \frac{p^a_H\left(s^a_L\right)}{p^a_L\left(s^a_L\right)} = \frac{s'_L}{1-s'_L} = \frac{s_L-\delta}{1-(s_L-\delta)}$$

I illustrate this construction in Figure 1.5.2d.

A buyer who receives an initial high signal s_H observes no additional information. She trades with the seller. However, a buyer who receives an initial low signal s_L receives additional information through $\mathcal{E}^a$. Absent this additional information, she would have rejected the seller. But a high signal s^a_H from $\mathcal{E}^a$ "positively overrides" her initial verdict: now, she trades.

So, stronger bad news encourages trade: a seller who previously would have been rejected by every buyer might now trade with one. The effect this has on total surplus is less clear. What can we infer about the quality of this seller?

The seller who would have been rejected by every buyer absent the additional information $\mathcal{E}^a$ supplies; so, initially all buyers observed low signals, s_L. Upon the additional information conveyed by $\mathcal{E}^a$, some buyer revises her decision to an acceptance; so, at least one buyer must have observed a high signal s_H^a from this additional experiment, but the rest must have observed the low signal s_L^a. However, the rest observed the low signal s_L^a from this additional experiment. Whether the expected surplus from trade is positive depends on *how many* did so:

$$\mathbb{P}\left(\theta = H \mid \geq 1 \text{ buyer observed } s_H^a\right) = \sum_{k=1}^{n} \left[\begin{array}{l} \mathbb{P}\left(\theta = H \mid k \text{ buyers observed } s_H^a\right) \\ \times \mathbb{P}\left(k \text{ buyers observed } s_H^a \mid \geq 1 \text{ buyer obs. } s_H^a\right) \end{array} \right]$$

However, we may deduce that almost surely only one buyer observed the high signal s_H^a. To see this, observe the likelihood ratios of the signals a buyer may observe from the experiment $\mathcal{E}^a$:

$$\frac{s_L^a}{1-s_L^a} = \frac{\frac{s_L-\delta}{1-(s_L-\delta)}}{\frac{s_L}{1-s_L}} \qquad\qquad \frac{s_H^a}{1-s_H^a} = \frac{\frac{s_H}{1-s_H}}{\frac{s_L}{1-s_L}}$$

While the likelihood ratio for the high signal s_H^a is constant, that for the low signal s_L^a converges to 1 as $\delta \downarrow 0$ and the improvement in buyers' information becomes "smaller". Due to the martingale property of likelihood ratios, these likelihood ratios must average to 1; so, the probability that buyers will observe the additional high signal s_H^a must vanish as $\delta \downarrow 0$.

So, the expected surplus from this trade is non-negative if and only if:

$$\lim_{\delta \downarrow 0} \mathbb{P}\left(\theta = H \mid \geq 1 \text{ buyer observed } s_H^a\right) - c = \mathbb{P}\left(\theta = H \mid 1 \text{ buyer observed } s_H^a\right) - c \geq 0$$
$$\iff \frac{\rho}{1-\rho} \times \left[\frac{s_L}{1-s_L}\right]^{n-1} \times \frac{s_H}{1-s_H} \geq \frac{c}{1-c}$$

Marginally stronger bad news allows the most adversely selected seller to trade. Had buyers' not received additional information, he would not have traded. Because they do,

one buyer accepts him. The RHS of the expression above reflects this: the expected surplus from trading with this seller is positive if and only if the high signal s_H^a observed by that one buyer overpowers the low signals s_L^a observed by the remaining $n-1$ buyers.[33]

This is a stark condition; when it holds, we may say that *adverse selection is irrelevant*: a buyer who observes the high signal need not be concerned about any previous rejections the seller may have received. It is also closely linked to the threshold Theorem 1.2 identifies, which Proposition 1.4 characterises.

Proposition 1.4. Let buyers' experiment be binary. Then, equilibrium* total surplus weakly decreases with stronger bad news (lower s_L) when:

$$\frac{\rho}{1-\rho} \times \max\left\{\frac{s_L}{1-s_L}, \left[\frac{s_L}{1-s_L}\right]^{n-1} \times \frac{s_H}{1-s_H}\right\} \leq \frac{c}{1-c}$$

This condition is also necessary in the least selective equilibrium.

Corollary 1.5. Where buyers' experiment is binary and $\rho \leq c$, equilibrium* total surplus weakly decreases with stronger bad news (lower s_L) when $\left(\frac{s_L}{1-s_L}\right)^{n-1} \times \frac{s_H}{1-s_H} \leq \frac{c}{1-c}$.

Proposition 1.4 reveals that total surplus falls with stronger bad news once bad news is strong enough to:

- violate the condition we informally identified as the "irrelevance of adverse selection", and
- for buyers to reject the seller with positive probability in their equilibrium* strategies.

Under some parameter constellations, even when adverse selection is no longer irrelevant, bad news might need to get stronger before, in equilibrium, buyers start rejecting the seller upon a low signal. Before the strength of bad news hits this critical mark, total surplus equals the no-information benchmark. Once it does, total surplus experiences a one time upward jump. Thereafter, stronger bad news decreases total surplus. Corollary 1.5 reveals, however, that this is not a concern when $\rho < c$: then, buyers always reject upon a low signal, so stronger bad news decreases total surplus whenever the condition we dubbed "the irrelevance of adverse selection" is satisfied.

[33] Note that $s_L^a \to s_L$ as $\delta \downarrow 0$.

1.5.2 Finite Experiments

Here, I show how the ideas we through Theorem 1.2 generalise to Blackwell improvements of experiments with an arbitrary finite number of outcomes. We cannot deploy those ideas immediately: first, for non-binary experiments, the ideas of stronger good and bad news lose meaning; second, Blackwell improvements of such experiments are complex—they cannot be described by simple changes in signals' likelihood ratios. Nonetheless, the core idea behind Theorem 1.2 supplies the answer: whether total surplus improves depends on whether an improvement is a *positive override*—information about a seller who would be rejected—or a *negative override*—information about a seller who would be approved.

Before I state Theorem 1.3, I establish two definitions that will provide useful langugage: Definition 1.2 formalises the notion of a positive and negative override, and Definition 1.5 formalises the notion of the "irrelevance of adverse selection".

Definition 1.2. Enumerate the joint outcome set of the experiments $\mathcal{E} = (\mathbf{S}, p_L, p_H)$ and $\mathcal{E}' = (\mathbf{S}', p'_L, p'_H)$ as $\mathbf{S} \cup \mathbf{S}' = \{s_1, s_2, ..., s_M\}$. The experiment $\mathcal{E}'$ differs from $\mathcal{E}$ by a *local mean preserving spread* (or, *local spread*) *at* $s_j \in \mathbf{S}$ if:

$$p'_\theta(s_j) = 0 \qquad p_\theta(s_{j+1}) = p_\theta(s_{j-1}) = 0 \qquad p'_\theta(s_{j+1}) + p'_\theta(s_{j-1}) = p_\theta(s_j)$$

and $p'_\theta(s) = p_\theta(s)$ for any $s \in \mathbf{S}' \cup \mathbf{S} \setminus \{s_{j-1}, s_j, s_{j+1}\}$.

A local spread moves all the probability mass experiment $\mathcal{E}$ places on some signal s_j to two new signals—one better news about the asset's quality, s_{j+1}, one worse, s_{j-1}. In this sense, we can think of it as providing additional information to a buyer who observes the original signal $s_j \in \mathbf{S}$.

Every local spread is an ordinary mean preserving spread.[34] The converse is not true; local spreads must move *all* the probability mass on a signal, and they must move it to its neighbouring signals: the mass on s_j is spread to s_{j+1} and s_{j-1}. I visualise the construction of a local spread in Figure 1.5.3.[35]

[34] Specifically, a "3-part MPS", in the language of Rasmusen and Petrakis (1992). Mean preserving spreads are originally due to Muirhead (1900) and were popularised in economics by Rothschild and Stiglitz (1970).

[35] For clarity, I fix $\psi = 0.5$ in this illustration.

Restricting to local spreads is without loss for finite experiments[36]—every Blackwell improvement, and *a fortiori*, ordinary mean preserving spread can be constructed through a finite number of local spreads.

Remark 1.3. [Müller and Stoyan (n.d.), Theorem 1.5.29] An experiment $\mathcal{E}'$ is Blackwell more informative than another, $\mathcal{E}$, if and only if there is a finite sequence of experiments $\mathcal{E}_1, \mathcal{E}_2, ..., \mathcal{E}_k$ such that $\mathcal{E}_1 = \mathcal{E}$, $\mathcal{E}_k = \mathcal{E}'$, and $\mathcal{E}_{i+1}$ differs from $\mathcal{E}_i$ by a local spread.

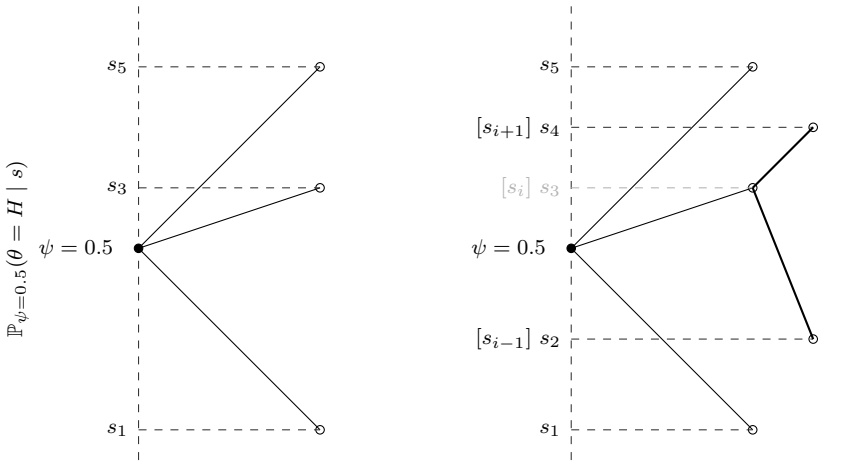

(a) Before local spread (b) After local spread

Figure 1.5.3: A Local Mean Preserving Spread

Definition 1.4. Let experiment $\mathcal{E}'$ differ from $\mathcal{E}$ by a local spread at s_j and $\sigma^\star_{\mathcal{E}}$ denote the equilibrium* strategy under experiment $\mathcal{E}$. This local spread is a *negative override* under equilibrium* if $\sigma^\star_{\mathcal{E}}(s_j) = 1$, and a *positive override* if $\sigma^\star_{\mathcal{E}}(s_j) = 0$.

Both positive and negative overrides are local spreads of a buyer's experiment, but they differ in *which* seller they inform the buyer about. A negative override is a local spread of a signal upon which a seller would have traded with the buyer. A positive override is a local spread of a signal upon which the seller would have been rejected by the buyer.

Definition 1.5. Let σ be a monotone strategy for a fixed experiment $\mathcal{E}$. *Adverse selection is σ-irrelevant for signal $s \in \mathbf{S}$* if:

$$\frac{\rho}{1-\rho} \times \left[\frac{r_H(\sigma;\mathcal{E})}{r_L(\sigma;\mathcal{E})}\right]^{n-1} \times \frac{s}{1-s} \geq \frac{c}{1-c}$$

[36] However, local mean preserving spreads are only defined for finite experiments; see Müller and Scarsini (2001) and Müller and Stoyan (n.d.).

When adverse selection is σ-irrelevant for a signal $s \in \mathbf{S}$, a buyer finds it optimal to trade upon the signal $s \in \mathbf{S}$ even if all other buyers rejected the seller—provided those buyers used the strategies σ.

Theorem 1.3. Let the experiment $\mathcal{E}'$ differ from $\mathcal{E}$ by a local spread at $s_j \in \mathbf{S}$. Equilibrium* total surplus is:

1. weakly greater under $\mathcal{E}'$ if the local spread is a negative override under equilibrium*.

2. weakly less under $\mathcal{E}'$ if the local spread is a positive override under equilibrium*, unless adverse selection is $\sigma^\star_{\mathcal{E}}$-irrelevant for signal s_{j+1}.

Theorem 1.3 shows that the effect of a local spread on total surplus depends on the *kind* of the spread. Negative overrides always increase total surplus. Positive overrides decrease it—unless adverse selection is irrelevant for a buyer who receives the override.

To prove Theorem 1.3, I show that a negative override indeed pushes buyers to reject the seller more often in the new equilibrium*. This always increases total surplus. A positive override pushes buyers to trade more often with him. This decreases total surplus unless adverse selection is $\sigma^\star_{\mathcal{E}}$-irrelevant for signal s_{j+1}—the new signal upon which the buyer may approve the seller.

This exercise is complicated by the fact that identifying how buyers' interim beliefs change when their experiment changes is infeasible beyond the simplest cases. Studying "local", not ordinary, mean preserving spreads is crucial for tractability; this allows me to identify how equilibria respond to improvements in information without needing to pinpoint changes in interim beliefs.

Unlike Theorem 1.2, an analyst who wishes to use Theorem 1.3 to identify the effect of an improvement on total surplus needs knowledge of buyers' equilibrium* strategies. In practice, the analyst might want to remain agnostic about equilibrium* strategies. To alleviate this concern, Proposition 1.6 offers a sufficient condition for a positive override to decrease total surplus in the most selective equilibrium.

Proposition 1.6. Let the experiment $\mathcal{E}'$ differ from $\mathcal{E}$ by a local spread at s_j. Total surplus in the most selective equilibrium is lower under $\mathcal{E}'$ if the following conditions hold:

$$\frac{\rho}{1-\rho} \times \left(\frac{s_j}{1-s_j}\right) \leq \frac{c}{1-c} \qquad \text{and} \qquad \frac{\rho}{1-\rho} \times \left(\frac{s_j}{1-s_j}\right)^{n-1} \times \frac{s_{j+1}}{1-s_{j+1}} \leq \frac{c}{1-c}$$

The sufficient conditions in Proposition 1.6 strengthen the necessary and sufficient conditions supplied by Theorem 1.3. The condition on the left ensures that the local spread is a negative override: a buyer rejects the seller upon s_j in any equilibrium since interim beliefs always lie below the prior belief, ρ. The condition on the right strengthens the irrelevance condition for adverse selection: it requires that a rejection be optimal even if the $n-1$ rejections the seller received were due to the best signals below s_{j+1}, s_j.

1.6 Maximising Surplus by Coarsening Information

Section 1.5 showed that efficiency might be lower in a market where buyers are better informed: more information in the market accentuates the adverse selection problem each buyer faces, counteracting each buyer's improved ability to screen the asset's quality. This invites a question: if a regulator could coarsen buyers' information—perhaps through banning the use and dissemination of certain data—could this raise efficiency? How should a regulator who can use this tool to maximise total surplus in the market go about this exercise?

In this section, I consider the problem of a regulator who wishes to garble buyers' experiment $\mathcal{E}$ in order to maximise total surplus in (the most selective) equilibrium. As in the previous section, I take the number of buyers n to be a primitive. The regulator can choose any finite garbling $\mathcal{E}^G = \left(\mathbf{S}^G, p_L^G, p_H^G\right)$ of buyers' experiment $\mathcal{E}$; i.e. any finite set of outcomes $\mathbf{S}^G = \left\{s_1^G, s_2^G, ..., s_R^G\right\}$ and conditional distributions $p_\theta^G(.)$ over it such that for some Markov matrix $\mathbf{T}_{m\times R}$:

$$\underbrace{\begin{bmatrix} p_L(s_1) & \cdots & p_L(s_m) \\ p_H(s_1) & \cdots & p_H(s_m) \end{bmatrix}}_{=\mathbf{P}} \times \mathbf{T} = \underbrace{\begin{bmatrix} p_L^G(s_1^G) & \cdots & p_L^G(s_R^G) \\ p_H^G(s_1^G) & \cdots & p_H^G(s_R^G) \end{bmatrix}}_{=\mathbf{P^G}}$$

We can interpret this as a coarsening of the original data available to each buyer. This

original data is generated with the process that matrix $\mathbf{P}$ describes. The regulator garbles this data into a set of summary statistics through the process described by the Markov matrix $\mathbf{T}$. The buyer only observes this set of summary statistics, whose data generating process conditional on the asset's quality can now described by the matrix $\mathbf{P^G}$.

Once the regulator chooses the garbled experiment $\mathcal{E}^G$, the game proceeds as before; only, buyers' experiment $\mathcal{E}$ is replaced by $\mathcal{E}^G$. An equilibrium, as before, is a pair $\left(\sigma^G, \psi^G\right)$ such that the strategy $\sigma^G : \mathbf{S}^G \to [0,1]$ is optimal given the interim belief ψ^G, and the interim belief ψ^G is consistent with the strategy σ^G. I call this the game induced by the garbling $\mathcal{E}^G$. I call a garbling $\mathcal{E}^G$ regulator-optimal if total surplus in the most selective equilibrium of the game it induces weakly exceeds equilibrium total surplus in the game induced by any other garbling of $\mathcal{E}$.

A particularly simple class of garblings are the *monotone binary garblings*, which (i) have two possible outcomes, $|\mathbf{S}^G| = 2$, and (ii) for a cutoff signal $s^* \in \mathbf{S}$, the entries $\{t_{ij}\}$ of matrix $\mathbf{T}_{m\times 2}$ for which $\mathbf{P} \times \mathbf{T} = \mathbf{P^G}$ are given by:

$$t_{i1} = \begin{cases} 1 & i < i^* \\ \in [0,1] & i = i^* \\ 0 & i > i^* \end{cases} \qquad\qquad t_{i2} = 1 - t_{i1}$$

I refer to the signal $s_{i^*} \in \mathbf{S}$ as the *threshold signal* of the garbling $\mathcal{E}^G$.

A monotone binary garbling gives the buyer an "acceptance recommendation", s_H^G, when her original signal realises above a threshold signal $s_{i^*} \in \mathbf{S}$, and a "rejection recommendation", s_L^G, whenever it lies below it. Following these recommendations—accepting trade upon the signal s_H^G and rejecting when it upon the signal s_L^G—need not be an equilibrium strategy in the game induced by the coarsened experiment $\mathcal{E}^G$. When it does, I say that the garbling $\mathcal{E}^G$ *produces incentive compatible recommendations.*

Definition 1.6. A monotone binary garbling $\mathcal{E}^G$ *produces incentive compatible (IC) recommendations* if the strategy σ^G, defined below, is an equilibrium strategy in the game

induced by $\mathcal{E}^G$:

$$\sigma^G(s^G) := \begin{cases} 0 & s^G = s_L^G \\ 1 & s^G = s_H^G \end{cases}$$

Lemma 1.7. Where it exists, the regulator-optimal garbling is monotone binary and produces IC recommendations.

The reason that the regulator can restrict herself to binary garblings is closely connected to a fundamental principle in information design. Ultimately, a buyer distils the information relayed by the garbled experiment into which action it recommends: an acceptance, or a rejection. The regulator can distil information herself—revealing only a recommended action to a buyer.[37] The regulator wishes trade to be as likely as possible with a High quality seller but not with a Low quality one, so she prefers monotone recommendations—those which recommend an acceptance above a threshold signal. Lemma 1.7 establishes that monotone recommendations are not at the expense of incentive compatibility: every garbling that supplies IC recommendations is outperformed by a monotone binary garbling that supplies IC recommendations.

We can adopt the selectivity order for strategies to a selectivity order for monotone binary garblings, too.

Definition 1.7. A monotone garbling $\mathcal{E}^G$ of $\mathcal{E}$ is *more selective* than another, $\mathcal{E}^{G'}$, if $p_\theta^G\left(s_H^G\right) \leq p_\theta^{G'}\left(s_H^{G'}\right)$ for all $\theta \in \{L, H\}$.

Like monotone strategies, and for the same reason, selectivity is a complete order over the set of monotone binary garblings. Similarly, we can adapt the "irrelevance of adverse selection" notion to monotone binary garblings as well.

Definition 1.8. Let $\mathcal{E}^G$ be a monotone binary garbling of $\mathcal{E}$, with the threshold signal $s^* \in \mathbf{S}$. *Adverse selection is irrelevant* under $\mathcal{E}^G$ either if:

$$\frac{\rho}{1-\rho} \times \frac{p_H(s^*)}{p_L(s^*)} \times \left(\frac{p_H^G(s_L^G)}{p_L^G(s_L^G)}\right)^{n-1} \geq \frac{c}{1-c}$$

[37] Note that coarsening buyers' experiment $\mathcal{E}$ may also create new equilibrium outcomes, some of which yield lower payoff than the previous least selective equilibrium. Our focus on the most selective equilibrium—besides being the appropriate focus for this exercise—frees us from the need to worry about this complication and utilise this fundamental principle.

or either of the following two conditions hold:

1. $\mathcal{E}^G$ recommends no acceptances; i.e. $p_\theta(s_H^G) = 0$.

2. $\mathcal{E}^G$ recommends no rejections; i.e. $p_\theta(s_L^G) = 0$ and $\frac{\rho}{1-\rho} \times \left(\frac{s_1}{1-s_1}\right)^n \geq \frac{c}{1-c}$.

Adverse selection is irrelevant under a garbling if a buyer who receives an "acceptance" recommendation need not be concerned about the number of buyers who received "rejection" recommendations—even if she believes that her acceptance recommendation is based on the lowest signal that might have triggered it. This condition is also satisfied if the garbling never recommends an "acceptance", or if, even if all n buyers were to observe the lowest signal under experiment $\mathcal{E}$, the expected gains from trade would be positive. When this condition is violated, I say that adverse selection is *not* irrelevant under $\mathcal{E}^G$.

Proposition 1.8. If the least selective monotone binary garbling under which adverse selection is irrelevant produces IC recommendations, it is the regulator-optimal garbling. Otherwise, the regulator-optimal garbling is either:

- the least selective garbling under which adverse selection is irrelevant, or
- the most selective garbling under which adverse selection is not irrelevant

among monotone binary garblings which produce IC recommendations.

Corollary 1.9. When the seller's reservation value c weakly exceeds the prior belief ρ, the regulator-optimal garbling is the least selective monotone binary garbling under which adverse selection is irrelevant.

Proposition 1.8 shows that the regulator wishes to recommend a rejection following every signal a buyer could observe, unless adverse selection is irrelevant at that signal. Although the regulator wishes to maximise a buyer's *expected* contribution to trade surplus, she focuses on the "worst case" where a buyer is the last to receive the seller.

The intuition is intimately connected to the insight the previous section delivers: unless "adverse selection is irrelevant", information which pushes buyers to accept trade more often can harm efficiency. The regulator wishes to censor such information by coarsening buyers' experiment: if she were not bound by buyers' incentives to follow her recommendations, her optimal garbling would bundle every outcome of the original experiment $\mathcal{E}$ into a "rejection

recommendation" unless adverse selection is irrelevant at that outcome. Corollary 1.9 establishes that when the seller's reservation value c weakly exceeds buyers' prior belief ρ, such recommendations are IC, and hence are adopted by the regulator.

1.7 Ultimatum Price Offers by Buyers

In this section, I relax the assumption that the seller trades with the first buyer willing to pay his reservation value, c. Instead, each buyer he visits offers the seller a take-it-or-leave-it (ultimatum) price. The seller then decides whether to trade at this price or move on to the next buyer in search of a better offer. I show that in this extended game, there exists an equilibrium where a buyer never offers a price above the seller's value c. Furthermore and crucially, any level of total surplus that can be achieved in some equilibrium can be achieved in an equilibrium where a buyer never offers a price above the seller's value c. Thus, restricting to such equilibria is without loss when studying the possible levels of equilibrium total surplus.

As before, the seller visits the n prospective buyers sequentially and in a uniformly random order, and each buyer holds a private signal about the asset's quality. The seller holds a private signal about the asset's quality—this signal is the outcome of a Blackwell experiment $\dot{\mathcal{E}}$. Conditional on the asset's quality, the seller's and buyers' signals are mutually independent. The seller enjoys a surplus of $o - c$ if he trades with some buyer at price o, and a surplus of 0 if he does not trade. A buyer enjoys a surplus of $\mathbb{1}\{\theta = H\} - o$ if she trades with the seller at price o, and a surplus of 0 if she does not trade. So, trade generates a total surplus of $\mathbb{1}\{\theta = H\} - c$.

Each buyer he visits offers the seller a take-it-or-leave-it price for the asset. A buyer's strategy $\omega : \mathbf{S} \to \Delta(\{0\} \cup [c, 1])$ maps every signal she might observe to a distribution from which her offer is drawn.[38]

Once the seller receives an offer of $o_k \in [0, 1]$ from the kth buyer he visits, he updates the probability q_{k-1} he places on the asset having High quality to a probability q_k. He does so with a belief updating rule described by $q_k = \zeta_\omega(q_{k-1}, o_k)$, where $\zeta_\omega(q_{k-1}, o_k) : [0, 1]^2 \to [0, 1]$ is a measurable function that is continuous in its first argu-

[38] I exclude offers in the interval $(0, c)$ without loss of generality: such offers will always be rejected by the seller since they are below his reservation value; the buyer might as well offer 0 instead.

ment and can depend on the strategy ω the seller believes buyers to be using.[39] This belief updating procedure is commonly known to all players.

The seller then decides whether to take or refuse the buyer's offer; his strategy $\chi_k(q_k, o_k) : [0,1]^2 \to [0,1]$ is a measurable mapping from his belief that the asset has High quality, q_k, and the kth buyer's offer, o_k, to a probability that he takes this k^{th} offer.

Given a strategy ω^* for buyers and the seller's belief q_k that the asset has High quality, denote the maximum expected surplus the seller can achieve after he rejects his k^{th} offer as $V_k^*(q_k)$. The seller secures zero surplus when he does not trade with any buyer; so, $V_n^*(q_n) = 0$ for all $q_n \in [0,1]$. We can then calculate each $V_k^*(q_k)$ recursively:

$$V_k^*(q_k) = \sum_{s \in \mathbf{S}} \left(q_k \times p_H(s) + (1 - q_k) \times p_L(s)\right) \times \int_0^1 \max\left\{V_{k+1}^*\left(\zeta_\omega(q_k, o_{k+1})\right), o_{k+1} - c\right\} d\omega^*(s)\left(o_{k+1}\right)$$

Given a strategy ω^* for buyers and $\{\chi_k^*\}_{k=1}^n$ for the seller, each buyer uses Bayes' Rule to calculate the probability that (i) she is the kth buyer to be visited given the seller has not traded yet, (ii) given the seller is in his k^{th} visit, he will take an offer of o_k, and (iii) given the seller takes that offer in his k^{th} visit and the private signal $s \in \mathbf{S}$ the buyer observed, the quality is High. Denote these probabilities as $\kappa^*(k)$, τ_{k,o_k}^*, and $\mathcal{H}_{k,o_k,s}^*$.

The strategies ω^* and $\{\chi_k^*\}_{k=1}^n$ form an equilibrium of this extended model if and only if:

1. A buyer's offer maximises her expected surplus given her beliefs about the quality and behaviour of the seller:

$$\text{supp } \omega^*(s) \subseteq \arg\max_{o \in [0,1]} \sum_{k=1}^{n} \kappa^*(k) \times \tau_{k,o}^* \times \left[\mathcal{H}_{k,o,s}^* - c\right]$$

2. The seller takes an offer if and only if his surplus from doing so weakly exceeds the expected surplus from refusing it and continuing his visits:

$$\chi_k^*(q_k, o_k) = \mathbb{1}\left\{o_k - c \geq V_k^*(q_k)\right\}$$

[39] For instance, this updating rule might restrict the seller to update using Bayes' Rule wherever possible.

Proposition 1.10. There exists an equilibrium where the seller trades whenever he is offered any price above c, and a buyer offers the seller either a price of c, or 0. Furthermore, any level of total surplus that can be achieved in equilibrium can be achieved by one such equilibrium.

Proof. The first part of the Proposition is seen easily: unless no buyer ever offers a price above c, no buyer has an incentive to offer such a price—the seller will accept the first offer weakly above c. By Proposition 1.1, then, the existence of such an equilibrium is guaranteed.

I prove the second part in two steps. I first show that buyers can never be offering a price above c in equilibrium, unless a seller of some quality receives the same price offer from every buyer. Then, I show that if a seller of some quality does receive the same price offer from every buyer, the total surplus in this equilibrium equals total surplus in another equilibrium where the seller is never offered a price above c.

Step 1: Observe that the seller's value function $V_k^*(.)$ is continuous: $V_k^*(.)$ is continuous provided $V_{k+1}^*(.)$ is, and $V_n^*(.)$ is a constant function. So, choose $k^{\max} \in \{1, 2, ..., n\}$ and $q^{\max} \in [0, 1]$ such that $V_{k^{\max}}^*(q^{\max}) = V^{\max} := \max_{k \in \{1,2,...,n\}, q \in [0,1]} V_k^*(q)$. If $V^{\max} = 0$, a seller at most expects an offer of c; we are done. So, let $V^{\max} > 0$. Then, no buyer can be offering a price above $V^{\max} + c$ in equilibrium: any seller accepts such a price, so the buyer would have a profitable deviation to $V^{\max} + c$. But then, a seller with a belief $q^{\max}$ must be expecting an offer of $V^{\max} + c$ with certainty in his $k + 1^{\text{th}}$ visit. This implies that for some quality $\theta \in \{L, H\}$, buyers offer the price $V^{\max} + c$ after any signal $s \in \text{supp } p_\theta$.

Step 2: Consider an equilibrium where for some quality $\theta \in \{L, H\}$, buyers offer the price $V^{\max} + c > c$ after any signal $s \in \mathbf{S}$ such that $s \in \text{supp } p_\theta$.

1. If buyers offer this price for any $s \in \mathbf{S}$, there is another equilibrium where they instead offer the seller c after every signal. In both equilibria, the seller trades with the first buyer he visits, regardless of his quality.

2. If buyers offer this price for any $s \in \text{supp } p_H$, the seller trades with the first buyer who offers this price whenever he has High quality. Buyers must be offering 0 after any $s \in \text{supp } p_L \setminus \text{supp } p_H$. There is then another equilibrium achieving the same

total surplus where buyers offer c after any $s \in \text{supp } p_H$ but 0 after any any $s \in \text{supp } p_L \setminus \text{supp } p_H$. The seller trades with the first buyer who offers c.

3. If buyers offer this price for any $s \in \text{supp } p_L$, the seller trades with the first buyer whenever he has Low quality. Buyers cannot be ever offering the seller a price of 0 following $s \in \text{supp } p_H \setminus \text{supp } p_L$; otherwise, a buyer would have a strictly positive deviation to an offer of c, which guarantees a profitable trade when she is the last buyer to be visited by the seller. So, the seller trades with probability one regardless of his quality. There is then another equilibrium where buyers offer the seller a price of c regardless of their signal, and the seller trades with probability one at this price. This equilibrium achieves the same total surplus.

□

1.8 Supplemental Appendix

Section 1.4 claimed that the outcome in a large market might be partially informative about the common value of the asset. This section presents a numerical example to demonstrate this claim.

Let buyers' experiment $\mathcal{E}$ be binary; $\mathbf{S} = \{0.2, 0.8\}$ and:

$$p_L(s) = \begin{cases} 0.8 & s = 0.2 \\ 0.2 & s = 0.8 \end{cases} \qquad p_H(s) = \begin{cases} 0.2 & s = 0.2 \\ 0.8 & s = 0.8 \end{cases}$$

Furthermore, let buyers' common prior be $\rho = 0.5$ and the seller's reservation value be $c = 0.6$. For any number of buyers, the equilibrium of this game is unique:[40] buyers always reject upon the low signal, $\sigma_n^*(0.2) = 0$, but accept with some probability upon the high signal, $\sigma_n^*(0.8) \in [0, 1]$.

Figures 1.8.b through 1.8.f plot (a) total surplus in the market, and the probabilities that (b) a buyer trades upon receiving the high signal, (c and d) some buyer trades with the seller given his quality, (e and f) the probability that a seller has High quality given he trades with some buyer or no buyer, as the number of buyers rises from 1 to 50.

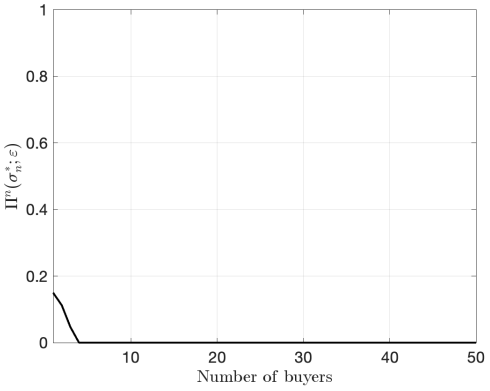

Figure 1.8.a: Total surplus

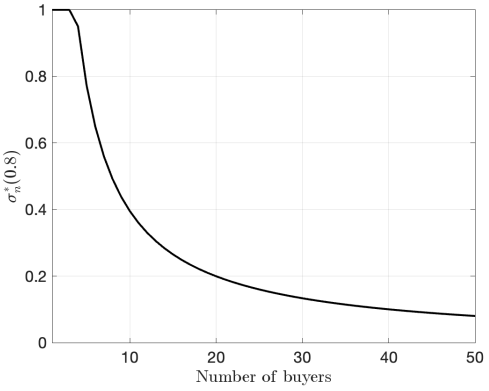

Figure 1.8.b: Probability that buyer accepts upon $s = 0.8$

[40] This follows a simple argument. No buyer can accept trade upon the low signal, since $\psi^* \leq 0.5$, so $\mathbb{P}_{\psi^*}(\theta = H \mid 0.2) \leq 0.2 < c$. As Lemma 1.15 shows, the interim belief is strictly decreasing in the probability that buyers approve upon the high signal. Thus, buyers either always accept trade at this signal, or there is a unique interior probability of acceptance that—given the interim belief consistent with those strategies—buyers are indifferent to accept.

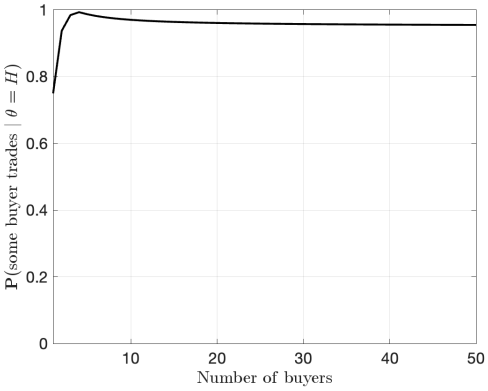

Figure 1.8.c: Probability that some buyer trades when $\theta = H$

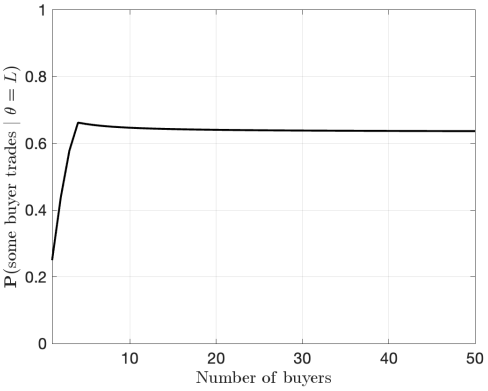

Figure 1.8.d: Probability that some buyer trades when $\theta = L$

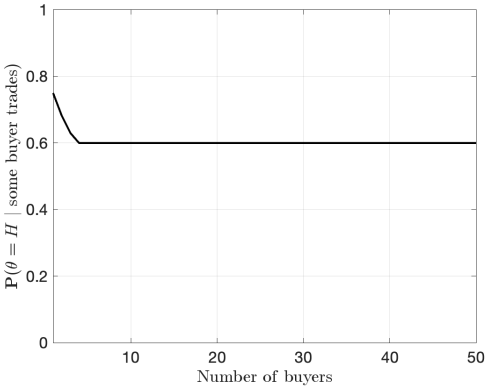

Figure 1.8.e: Probability that $\theta = H$ given some buyer trades

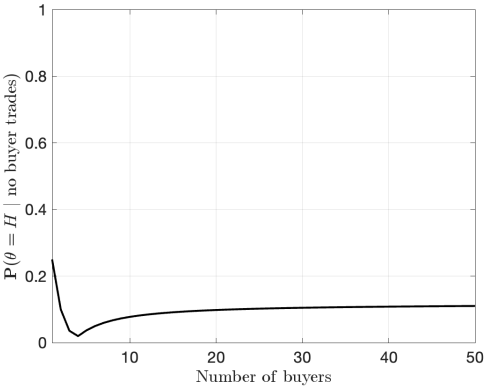

Figure 1.8.f: Probability that $\theta = H$ given no buyer trades

Once the number of buyers in the market exceeds 3, total surplus in the market equals 0—surplus in the no-information benchmark. A buyer only trades if she receives the high signal. Even then, she rejects the seller with a strictly positive probability. This probability increases with the number of buyers in the market. However, the probability that some buyer trades with the seller reaches a constant level: 0.95 if he has High quality, and 0.63 if he has Low quality. Despite having no bearing on market participants' surplus, trade is informative about the seller's quality: the probability that the seller has High quality conditional on trading is 0.6 not $\rho = 0.5$; and the probability she has High quality conditional on not trading reaches 0.11.[41]

[41]For clarity, the figures illustrate results as $n \to 50$; however, these asymptotic values remain valid as $n \to 1000$. As $n \to 1000$, the probability that a buyer accepts trade upon a high signal converges to 0.

1.9 Proof Appendix

1.9.1 Useful Definitions and Notation

In what follows, I occasionally operate with the likelihood ratios of beliefs for convenience. The reader can easily verify the identities:

$$\frac{\psi}{1-\psi} = \frac{\rho}{1-\rho} \times \frac{\nu_H(\sigma;\mathcal{E})}{\nu_L(\sigma;\mathcal{E})} \qquad \frac{\mathbb{P}_\psi(\theta = H \mid s_i)}{1-\mathbb{P}_\psi(\theta = H \mid s_i)} = \frac{\psi}{1-\psi} \times \frac{s_i}{1-s_i}$$

Through similar reasoning, the reader can verify that it is optimal for a buyer to accept trade when:

$$\frac{\mathbb{P}_\psi(\theta = H \mid s_i)}{1-\mathbb{P}_\psi(\theta = H \mid s_i)} > \frac{c}{1-c}$$

Some strategies require buyers to randomise upon observing a particular outcome. To facilitate technical discussion, where it is warranted I assume that each buyer observes the realisation of a *tie-breaking signal* $u \sim U[0,1]$ alongside the outcome of her experiment. This signal is not informative about the asset's quality: it is distributed independently from it conditional on the experiment's outcome. I denote the outcome of buyer i's experiment as s^i and her tie-breaking signal as u^i. Without loss, buyer i accepts trade if and only if $\sigma(s^i) \leq u^i$; where σ is her strategy. I call the pair (s^i, u^i) the *score* buyer i observes for the seller.

Definition 1.9. The tuple $Z^i = (s^i, u^i)$, where $u^i \overset{IID}{\sim} U[0,1]$ is the *score* buyer i observes for the seller. The seller's *score profile* $\mathbf{z}$ is the set of scores each buyer observes; $\mathbf{z} = \{(s^i, u^i)\}_{i=1}^n$. Analogously, the seller's *signal profile* $\mathbf{s} = \{s^i\}_{i=1}^n$ is the set of outcomes of each buyer's experiment.

Some proofs in Section 1.9.3 require comparing interim beliefs across pairs of strategies and experiments; $(\sigma, \mathcal{E})$. For convenience, I define the mapping from such a pair to the interim belief consistent with them as $\Psi(.;\mathcal{E}) : [0,1]^n \to [0,1]$:

$$\Psi(\sigma;\mathcal{E}) := \frac{\rho \times \nu_H(\sigma;\mathcal{E})}{\rho \times \nu_H(\sigma;\mathcal{E}) + (1-\rho) \times \nu_L(\sigma;\mathcal{E})}$$

Wherever necessary, I treat each strategy $\sigma : \mathbf{S} \to [0,1]$ for an experiment $\mathcal{E}$ as a vector

in the compact set $[0,1]^m \subset \mathbb{R}^n$. This is a finite dimensional vector space, so I endow it with the metric induced by the taxicab norm without loss of generality (see Kreyszig (1978) Theorem 2.4-5):

$$||\sigma' - \sigma|| = \sum_{j=1}^{m} |\sigma'(s_j) - \sigma(s_j)| \qquad \text{for any two strategies } \sigma' \text{ and } \sigma$$

Note that the interim belief function $\Psi(.;\mathcal{E})$ is thus a continuous function of buyers' strategies.[42]

Definition 1.10. Where the experiment $\mathcal{E}$ is binary, s_L^{mute} is the *strongest* level of bad news for which there is an equilibrium where a buyer trades regardless of her signal:

$$\frac{\rho}{1-\rho} \times \frac{s_L^{\text{mute}}}{1-s_L^{\text{mute}}} = \frac{c}{1-c}$$

1.9.2 Omitted Results

Lemma 1.11. Let σ^* and σ be two monotone strategies such that (i) σ^* is more selective than σ, and (ii) σ^* is an equilibrium strategy. Then, total surplus is higher under σ^*: $\Pi(\sigma^*;\mathcal{E}) \geq \Pi(\sigma;\mathcal{E})$.

Proof. Let $\mathbf{z}$ be the seller's *score profile*. Take an equilibrium strategy σ^* and a less selective strategy σ such that:

$$\sigma(s) - \sigma^*(s) = \begin{cases} \varepsilon & s = \underline{s} \\ 0 & \text{otherwise} \end{cases}$$

for some $\varepsilon > 0$, where $\underline{s} := \min\{s \in S : \sigma^*(s) < 1\}$. I show that:

$$\lim_{\varepsilon \to 0} \Pi(\sigma;\varepsilon) - \Pi(\sigma^*;\varepsilon) \leq 0$$

By Lemma 1.14, this establishes the result.

Now, let $Z \subset (S \times [0,1])^n$ be the set of score profiles under which some buyer trades

[42] r_θ is continuous in σ; so, both the nominator and denominator are strictly positive continuous in σ.

under σ, but all buyers reject the seller under σ^*:

$$\mathbf{z} \in Z \iff \begin{array}{ll} \sigma^*(s^i) > u^i & \text{for all } i \in \{1, 2, \ldots, n\}, \\ & \text{and} \\ \sigma(s^i) \leq u^i & \text{for some } i \in \{1, 2, \ldots, n\}. \end{array}$$

Furthermore, for a given score profile $\mathbf{z}$, let $\#$ be the number of buyers whose observed scores are such that $\sigma(s^i) \geq u^i > \sigma^*(s^i)$. These buyers would accept trade under the strategy σ, but not under σ^*.

The seller's eventual outcome differs between the strategy profiles σ and σ^* if and only if his score profile $\mathbf{z}$ lies in Z. Furthermore, his eventual outcome can only change from a rejection by all buyers in σ^* to an approval by some buyer in σ. Thus:

$$\begin{aligned} \Pi(\sigma; \mathcal{E}) - \Pi(\sigma^*; \mathcal{E}) &= [\mathbb{P}(\theta = H \mid \mathbf{z} \in Z) - c] \times \mathbb{P}(\mathbf{z} \in Z) \\ &\propto \mathbb{P}(\theta = H \mid \mathbf{z} \in Z) - c \end{aligned}$$

Focus therefore, on the probability that $\theta = H$ given the seller's signal profile lies in Z:

$$\mathbb{P}(\theta = H \mid \mathbf{z} \in Z) = \sum_{i=1}^{n} \mathbb{P}(\theta = H \mid \# = i) \times \frac{\mathbb{P}(\# = i)}{\mathbb{P}(\mathbf{z} \in Z)}$$

Now note:

$$\mathbb{P}(\# = i \mid \theta) = (p_\theta(\underline{s}))^i \times (1 - p_\theta(\underline{s}))^{n-i} \times \varepsilon^i$$

and thus $\mathbb{P}(\# = i) \propto \varepsilon^i$. Since $\mathbb{P}(\mathbf{z} \in A) = \sum_{i=1}^{n} \mathbb{P}(\# = i)$, we have $\lim_{\varepsilon \to 0} \frac{\mathbb{P}(\#=i)}{\mathbb{P}(\mathbf{z} \in A)} = 0$ for any $i > 1$. Thus:

$$\lim_{\varepsilon \to 0} \mathbb{P}(\theta = H \mid \mathbf{z} \in A) - \mathbb{P}(\theta = H \mid \# = 1) = 0$$

I conclude the proof by showing that $\mathbb{P}(\theta = H \mid \# = 1) \leq c$ as $\varepsilon \to 0$:

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{\mathbb{P}(\theta = H \mid \# = 1)}{\mathbb{P}(\theta = L \mid \# = 1)} &= \lim_{\varepsilon \to 0} \frac{\mathbb{P}(\theta = H)}{\mathbb{P}(\theta = L)} \times \frac{\mathbb{P}(\# = 1 \mid \theta = H)}{\mathbb{P}(\# = 1 \mid \theta = L)} \\
&= \lim_{\varepsilon \to 0} \frac{\mathbb{P}(\theta = H)}{\mathbb{P}(\theta = L)} \times \left(\frac{r_H(\sigma; \mathcal{E})}{r_L(\sigma; \mathcal{E})} \right)^{n-1} \times \frac{p_H(\underline{s})}{p_L(\underline{s})} \\
&= \frac{\mathbb{P}(\theta = H)}{\mathbb{P}(\theta = L)} \times \left(\frac{r_H(\sigma^*; \mathcal{E})}{r_L(\sigma^*; \mathcal{E})} \right)^{n-1} \times \frac{p_H(\underline{s})}{p_L(\underline{s})} \\
&\leq \frac{\psi^*}{1 - \psi^*} \times \frac{p_H(\underline{s})}{p_L(\underline{s})} \leq \frac{c}{1 - c}
\end{aligned}$$

where $\psi^* = \Psi(\sigma^*; \mathcal{E})$ is the interim belief consistent with σ^*. The penultimate inequality holds due to the straightforward fact that:

$$\frac{\psi^*}{1 - \psi^*} = \frac{\rho}{1 - \rho} \times \frac{1 + r_H^* + ... + (r_H^*)^{n-1}}{1 + r_L^* + ... + (r_L^*)^{n-1}} \leq \frac{\rho}{1 - \rho} \times \left(\frac{r_H^*}{r_L^*} \right)^{n-1}$$

where $r_\theta^* := r_\theta(\sigma^*; \mathcal{E})$. The last inequality is due to the fact that $\underline{s} \in S$ is optimally rejected under σ^*.

□

Lemma 1.12. Suppose there is a single buyer, $n = 1$. Equilibrium total surplus under experiment $\mathcal{E}'$ exceeds that under $\mathcal{E}$ *regardless of the seller's reservation value* $c \in [0, 1]$ *and buyer's prior belief* $\rho \in [0, 1]$ if and only if $\mathcal{E}'$ is (Blackwell) more informative than $\mathcal{E}$.

Proof. The sufficiency part of this Lemma follows from Blackwell's Theorem (Blackwell and Girshick (1954), Theorem 12.2.2). To show necessity, I fix an arbitrary prior belief ρ for the evaluator.

Let q_j be the posterior belief the buyer forms about the asset's quality upon observing the outcome $s_j \in \mathbf{S}$:

$$q_j = \frac{\rho \times s_j}{\rho \times s_j + (1 - \rho) \times (1 - s_j)}$$

Furthermore, let $F(.)$ and $F'(.)$ be the CDFs of the posterior beliefs $\mathcal{E}$ and $\mathcal{E}'$ induce,

respectively, for this prior belief ρ:

$$F(q) = (1-\rho) \times \sum_{s \in \mathbf{S}: s \leq q} p_L(s) + \rho \times \sum_{s \in \mathbf{S}: s \leq x} p_H(s)$$
$$F'(q) = (1-\rho) \times \sum_{s \in \mathbf{S}: s \leq q} p'_L(s) + \rho \times \sum_{s \in \mathbf{S}: s \leq x} p'_H(s)$$

Equilibrium total surplus (and the buyer's expected payoff) under $\mathcal{E}$ is given by:

$$\int_c^1 (q-c)dF(q) = \int_c^1 qdF(q) - c \times (1 - F(c)) = (1-c) - \int_c^1 F(q)dq$$

An analogous expression gives equilibrium total surplus under $\mathcal{E}'$. For the former to exceed the latter for any $c \in [0,1]$, we must have:

$$\int_c^1 \left(F(q) - F'(q)\right) dq \geq 0$$

which is equivalent to $\mathcal{E}'$ being Blackwell more informative than $\mathcal{E}$.[43]

□

Lemma 1.13 proves useful when proving Theorem 1.3, the main result of Section 1.5.2. This Lemma can also be used towards an alternative and direct proof for the equilibrium existence claim of Proposition 1.1.

Lemma 1.13. For each $j \in \{1, 2, ..., m\}$, let σ_j be the strategy defined as:

$$\sigma_j(s) = \begin{cases} 0 & s < s_j \\ 1 & s \geq s_j \end{cases}$$

and ψ_j be the interim belief consistent with this strategy. Unless σ_j is itself an equilibrium strategy:

i There is an equilibrium strategy σ^* that is *more selective than* σ_j if $\mathbb{P}_{\psi_j}(\theta = H \mid s_j) < c$.

[43] See Müller and Stoyan (n.d.), Theorem 1.5.7. The Blackwell order between signal structures is equivalent to the convex order between the posterior belief distributions they induce; see Gentzkow and Kamenica (2016).

ii There is an equilibrium strategy σ^* that is *less selective than* σ_j if $\mathbb{P}_{\psi_j}(\theta = H \mid s_{j-1}) > c$.

Proof. Abusing notation slightly, I add two fully revealing outcomes s_0 and s_{m+1} to the set $\mathbf{S}$ (duplicating s_1 and s_m if either of them are already fully revealing), and denote the strategy which *never* accepts trade as σ_{m+1}:

$$\frac{s_{m+1}}{1-s_{m+1}} = \infty \qquad \frac{s_0}{1-s_0} = 0 \qquad \frac{\psi_{m+1}}{1-\psi_{m+1}} = \frac{\rho}{1-\rho}$$

Claim i.

The strategy σ_{m+1} is the most selective strategy buyers can adopt, and is an equilibrium strategy unless:

$$\frac{s_m}{1-s_m} \times \frac{\psi_{m+1}}{1-\psi_{m+1}} > \frac{c}{1-c}$$

So, assume this condition is satisfied. Likewise, the strategy σ_k for $k \geq j$ is an equilibrium if the following inequality is satisfied:

$$\frac{s_k}{1-s_k} \times \frac{\psi_k}{1-\psi_k} \geq \frac{c}{1-c} \geq \frac{s_{k-1}}{1-s_{k-1}} \times \frac{\psi_k}{1-\psi_k} \tag{1.9.1}$$

So, assume inequality 1.9.1 is violated for every $k \geq j$. This gives us:

$$\frac{s_{m+1}}{1-s_{m+1}} \times \frac{\psi_{m+1}}{1-\psi_{m+1}} > \frac{c}{1-c} > \frac{s_j}{1-s_j} \times \frac{\psi_j}{1-\psi_j} \tag{1.9.2}$$

where the last part of this inequality is by hypothesis.

Now, let $k^* \in \{j, j+1, ..., m\}$ be the first index for which the following inequality is satisfied:

$$\frac{s_{k^*+1}}{1-s_{k^*+1}} \times \frac{\psi_{k^*+1}}{1-\psi_{k^*+1}} \geq \frac{c}{1-c} \geq \frac{s_{k^*}}{1-s_{k^*}} \times \frac{\psi_{k^*}}{1-\psi_{k^*}}$$

such a k^* must exist due to inequality 1.9.2. But since inequality 1.9.1 is violated, we must have:

$$\frac{s_{k^*}}{1-s_{k^*}} \times \frac{\psi_{k^*+1}}{1-\psi_{k^*+1}} > \frac{c}{1-c} \geq \frac{s_{k^*}}{1-s_{k^*}} \times \frac{\psi_{k^*}}{1-\psi_{k^*}}$$

But since the function $\Psi(\sigma; \mathcal{E})$ is continuous in buyers' strategy σ, we can then find some

strategy σ^*:

$$\sigma^*(s) = \begin{cases} 1 & s > s_{k^*} \\ \in [0,1] & s = s_{k^*} \\ 0 & s < s_{k^*} \end{cases}$$

such that:

$$\frac{s_{k^*+1}}{1-s_{k^*+1}} \times \frac{\Psi(\sigma^*;\mathcal{E})}{1-\Psi(\sigma^*;\mathcal{E})} \geq \frac{c}{1-c} = \frac{s_{k^*}}{1-s_{k^*}} \times \frac{\Psi(\sigma^*;\mathcal{E})}{1-\Psi(\sigma^*;\mathcal{E})}$$

The strategy σ^* is thus an equilibrium strategy. It is clearly more selective than σ_j; since it is more selective than σ_{k^*}, where $k^* \geq j$.

Claim ii.

For any $k \in \{1, 2, ..., j\}$, the strategy σ_k is an equilibrium if the inequality 1.9.1 is satisfied. So, as earlier, assume 1.9.1 is violated for every such k. Then, we get:

$$\frac{s_j}{1-s_j} \times \frac{\psi_j}{1-\psi_j} \geq \frac{s_{j-1}}{1-s_{j-1}} \times \frac{\psi_j}{1-\psi_j} > \frac{c}{1-c} > \frac{s_1}{1-s_1} \times \frac{\psi_1}{1-\psi_1}$$

where the second inequality in the chain follows by hypothesis and the last inequality follows from the violation of inequality 1.9.1 for $k = 1$. We can now repeat the argument we constructed after inequality 1.9.2 to prove Claim i, to prove the existence of an equilibrium strategy σ^* that is less selective than σ_j.

□

1.9.3 Omitted Proofs

Proposition 1.1. Let Σ be the set of equilibrium strategies. Then:

1. Σ is non-empty and compact.
2. Any equilibrium strategy σ^* is monotone: for any $\sigma^* \in \Sigma$, $\sigma^*(s) > 0$ for some $s \in \mathbf{S}$ implies that $\sigma^*(s') = 1$ for every $s' \in \mathbf{S}'$ such that $s' > s$.
3. All equilibria exhibit adverse selection: $\psi^* \leq \rho$ for any interim belief ψ^* consistent with an equilibrium strategy $\sigma^* \in \Sigma$.

Proof. In what follows, I treat each strategy $\sigma : \mathbf{S} \to [0,1]$ as a vector in the compact set $[0,1]^m \subset \mathbb{R}^n$, endowed with the taxicab metric (see the end of Section 1.9.1). I start

by proving that any equilibrium strategy must be monotone and all equilibria exhibit adverse selection. Using these observations, I prove that the set of equilibrium strategies is non-empty and compact.

2. Any equilibrium strategy is monotone.

Any equilibrium strategy σ^* must be optimal against the interim belief ψ^* consistent with it. Whenever $\rho \in (0,1)$, $\psi^* = \Psi(\sigma^*; \mathcal{E}) \in (0,1)$, and so $\mathbb{P}_{\psi^*}(\theta = H \mid s') > \mathbb{P}_{\psi^*}(\theta = H \mid S = s)$ for $s', s \in \mathbf{S}$ such that $s' > s$.

3. All equilibria exhibit adverse selection.

A fortiori, $\Psi(\sigma; \mathcal{E}) \leq \rho$ for any monotone strategy σ. To see this, note that $p_H(.)$ first order stochastically dominates $p_L(.)$ since it likelihood ratio dominates it.[44] Therefore, $\nu_L(\sigma; \mathcal{E}) \geq \nu_H(\sigma; \mathcal{E})$. The result then follows since $\frac{\Psi(\sigma;\mathcal{E})}{1-\Psi(\sigma;\mathcal{E})} = \frac{\rho}{1-\rho} \times \frac{\nu_H(\sigma;\mathcal{E})}{\nu_L(\sigma;\mathcal{E})}$.

1. The set of equilibrium strategies is non-empty and compact.

i The set of equilibrium strategies is non-empty.

Define $\Phi(.) : [0,1]^m \to 2^{[0,1]^m}$ to be the buyers' *best response correspondence*. $\Phi(.)$ maps any strategy σ to the set of strategies that are optimal against the interim belief $\Psi(\sigma; \mathcal{E})$ it induces:

$$\Phi(\sigma) = \left\{\sigma' \in [0,1]^m : \sigma' \text{ is optimal against } \Psi(\sigma; \mathcal{E})\right\}$$

A strategy σ^* is an equilibrium strategy if and only if it is a fixed point of buyers' best response correspondence; $\sigma^* \in \Phi(\sigma^*)$. I establish that the correspondence Φ has at least such fixed point through Kakutani's Fixed Point Theorem.

Φ is trivially non-empty; every interim belief has some strategy optimal against it. It is also convex valued; if two distinct approval probabilities are optimal after some outcome $s \in \mathbf{S}$, *any* approval probability is optimal upon that outcome.

The only task that remains is to prove that Φ is upper-semi continuous. For this, take an arbitrary sequence of strategies $\{\sigma_n\}$ such that $\sigma_n \to \sigma_\infty$. Denote the interim beliefs consistent with these strategies as $\psi_n := \Psi(\sigma_n; \mathcal{E})$. Since $\Psi(.; \mathcal{E})$ is continuous in

[44] Theorem 1.C.1 in Shaked and Shanthikumar (2007).

buyers' strategies, we also have $\psi_n \to \psi_\infty$ where $\psi_\infty = \Psi(\sigma_\infty; \mathcal{E})$. Now, take a sequence of strategies $\{\sigma_n^*\}$ where $\sigma_n^* \in \Phi(\sigma_n)$. Note that every σ_n^* is monotone since optimality against any interim belief $\psi \in (0,1)$ requires monotonicity. We want to show that Φ is upper semicontinuous; i.e.:

$$\sigma_n^* \to \sigma_\infty^* \implies \sigma_\infty^* \in \Phi(\sigma_\infty)$$

By the Monotone Subsequence Theorem, the sequence $\{\sigma_n^*\}$ has a subsequence $\sigma_{n_k}^* \to \sigma_\infty^*$ of strategies whose norms $||\sigma_{n_k}^*||$ are monotone in their indices n_k. Here, I take the case where these norms are increasing, the proof is analogous for the opposite case. Since σ_∞^* is the limit of a subsequence of monotone strategies, it must be a monotone strategy too. Assuming otherwise leads to a contradiction; for any $s, s' \in \mathbf{S}$ such that $s' > s$:

$$\sigma_\infty^*(s) > 0 \;\&\; \sigma_\infty^*(s') < 1 \implies \exists N \in \mathbb{N} \text{ s.t. } \forall\, n_k \geq N \;\; \sigma_{n_k}^*(s) > 0 \;\&\; \sigma_{n_k}^*(s') < 1$$

Now let $\bar{s}$ be the highest outcome for which $\sigma_\infty^*(\bar{s}) > 0$. I show that:

- If $\sigma_\infty^*(\bar{s}) \in (0,1)$, then:
$$\frac{\psi_\infty}{1-\psi_\infty} \times \frac{\bar{s}}{1-\bar{s}} = \frac{c}{1-c}$$
- If $\sigma_\infty^*(\bar{s}) = 1$, then:
$$\frac{\psi_\infty}{1-\psi_\infty} \times \frac{s}{1-s} \begin{cases} \leq \frac{c}{1-c} & s < \bar{s} \\ \geq \frac{c}{1-c} & s \geq \bar{s} \end{cases}$$

The first case easily follows by noting that:

$$\begin{aligned} \sigma_\infty^*(\bar{s}) \in (0,1) &\implies \sigma_{n_k}^*(\bar{s}) \in (0,1) \\ &\implies \frac{\psi_{n_k}}{1-\psi_{n_k}} \times \frac{\bar{s}}{1-\bar{s}} = \frac{c}{1-c} \implies \frac{\psi_\infty}{1-\psi_\infty} \times \frac{\bar{s}}{1-\bar{s}} = \frac{c}{1-c} \end{aligned}$$

for all $n_k \geq N' \in \mathbb{N}$. The second case follows similarly, by noting that $\sigma_\infty^*(\bar{s}) = 1$ and $\sigma_\infty^*(s') = 0$ for all $s' < \bar{s}$ implies $\sigma_{n_k}^*(\bar{s}) > 0$ and $\sigma_{n_k}^*(s') = 0$ for all $n_k \geq N'' \in \mathbb{N}$.

ii The set of equilibrium strategies is compact.

Σ is a subset of $[0,1]^m$ and therefore bounded, hence it suffices to show that is closed. Let $\{\sigma_n^{**}\}$ be a sequence of equilibrium strategies. Note that this means $\sigma_n^{**} \in \Phi(\sigma_n^{**})$. Since $\Phi(.)$ is upper semicontinuous, $\sigma_n^{**} \to \sigma_\infty$ implies $\sigma_\infty \in \Phi(\sigma_\infty)$, and therefore σ_∞ is an equilibrium strategy itself.

□

Proposition 1.3. Equilibrium total surplus is bounded above by the full-information benchmark Π^f and below by the no-information benchmark $\Pi^\emptyset$. Furthermore, it is higher under more selective equilibrium strategies:

$$\max\{0, \rho - c\} = \Pi^\emptyset \leq \Pi(\sigma^{**}; \mathcal{E}) \leq \Pi(\sigma^*; \mathcal{E}) \leq \Pi^f = \rho \times (1 - c)$$

where σ^* and σ^{**} are two equilibrium strategies such that σ^{**} is more selective than σ^*.

Proof. This is an immediate corollary to Lemmas 1.11 and 1.14 below; both of independent interest. □

Lemma 1.14. Take three monotone strategies σ'', σ' and, σ, ordered from the least selective to the most. If $\Pi(\sigma'; \mathcal{E}) \leq \Pi(\sigma; \mathcal{E})$, then $\Pi(\sigma''; \mathcal{E}) \leq \Pi(\sigma'; \mathcal{E})$.

Proof. For the three strategies σ'', σ', and σ, consider three sets $Z, Z', Z'' \subset (S \times [0,1])^n$ where the seller's score profile $\mathbf{z}$ might lie:

$$\mathbf{z} \in \begin{cases} Z & \text{if } \mathbf{z} \text{ trades with some buyer under } \sigma'' \text{ but not } \sigma \\ Z' & \text{if } \mathbf{z} \text{ trades with some buyer under } \sigma' \text{ but not } \sigma \\ Z'' & \text{if } \mathbf{z} \text{ trades with some buyer under } \sigma'' \text{ but not } \sigma' \end{cases}$$

Notice that $Z' \cap Z'' = \emptyset$ and $Z' \cup Z'' = Z$. We can write the difference between total surplus under different strategies as:

$$\Pi(\sigma'; \mathcal{E}) - \Pi(\sigma; \mathcal{E}) = \mathbb{P}\left(\mathbf{z} \in Z'\right) \times \left[\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z'\right) - c\right]$$

and:

$$\Pi(\sigma''; \mathcal{E}) - \Pi(\sigma'; \mathcal{E}) = \mathbb{P}\left(\mathbf{z} \in Z''\right) \times \left[\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z''\right) - c\right]$$

Therefore we want to prove that:

$$\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z'\right) \leq c \implies \mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z''\right) \leq c$$

Now, note that $\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z\right)$ is a convex combination of $\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z'\right)$ and $\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z''\right)$. Furthermore:

$$\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z\right) \geq \mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z \cap Z''\right) = \mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z''\right)$$

which then implies:

$$\mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z''\right) \leq \mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z\right) \leq \mathbb{P}\left(\theta = H \mid \mathbf{z} \in Z'\right) \leq c$$

□

Theorem 1.1. Let $\Pi^n\left(\hat{\sigma}; \mathcal{E}\right)$ be total surplus under the most selective equilibrium in a market with n buyers. If $\mathcal{E}$ has an outcome that fully reveals High quality ($s_m = 1$), the sequence $\{\Pi^n\left(\hat{\sigma}; \mathcal{E}\right)\}_{n=1}^{\infty}$ is eventually increasing and converges to surplus in the full-information benchmark. Otherwise, it is eventually decreasing and converges to surplus in the no-information benchmark.

Proof. For each $j \in \{1, 2, ..., m\}$, define σ_j to be the strategy:

$$\sigma_j(s) := \begin{cases} 0 & s < s_j \\ 1 & s \geq s_j \end{cases}$$

Moreover, let $F_\theta(x) = \sum_{s \in \mathbf{S}: s \leq x} p_\theta(x)$. In a market with n buyers, the interim belief $\psi_{j:n}$ consistent with buyers using the strategies σ_j is then implicitly given by:

$$\frac{\psi_{j:n}}{1 - \psi_{j:n}} = \frac{\sum_{k=0}^{n-1} F_H(s_{j-1})^k}{\sum_{k=0}^{n-1} F_L(s_{j-1})^k}$$

Note that for all $j > 1$, the RHS is bounded and strictly decreasing in n, so the sequence $\{\psi_{j:n}\}$ is convergent.

Case 1: $s_m = 1$.

To prove the Theorem's statement for this case, I first show that $\psi_{m:n} \xrightarrow{n} 0$. Let X_n be the random variable that is uniformly distributed over the set $\left\{F_H(s_{m-1})^k\right\}_{k=0}^{n-1}$. Then, note that:

$$\frac{\psi_{m:n}}{1-\psi_{m:n}} = \frac{\sum_{k=0}^{n-1} F_H(s_{m-1})^k}{\sum_{k=0}^{n-1} F_L(s_{m-1})^k} = \mathbb{E}\left[X_n\right]$$

Now, fix any $x > 0$. Since $F_H(s_{m-1})^k$ is strictly decreasing in k, for any $\delta < 1$ of our choice, we can find some $N_{x;\delta} \in \mathbb{N}$ such that for all $n \geq N_{x;\delta}$ implies $\mathbb{P}\left(X_n \leq x\right) \geq \delta$ and $\mathbb{E}\left[X_n\right] \leq \delta x + (1-\delta)$. Fixing $x = \frac{\varepsilon}{2\delta}$ for some $\varepsilon > 0$ small, we have $\mathbb{E}\left[X_n\right] \leq \frac{\varepsilon}{2} + 1 - \delta$. Since we can take δ arbitrarily close to 1, this shows that $\mathbb{E}\left[X_n\right] \to 0$, proving this first claim.

Since buyers must always accept to trade upon observing $s_m = 1$ in equilibrium, this implies that there is some $N \in \mathbb{N}$ for which σ_m is the most selective equilibrium strategy for all $n \geq N$. So, for $n \geq N$, a seller with a Low quality asset never trades. Moreover, every additional buyer increases the probability that a seller with a High quality asset trades. As $n \to \infty$, such a seller trades almost surely. We thus prove that $\Pi^n\left(\hat{\sigma}; \mathcal{E}\right) \to \Pi^f$.

Case 2: $s_m < 1$

The case where buyers' experiment $\mathcal{E}$ is uninformative is trivial; it always yields the no-information benchmark. So, I assume that $s_m \neq s_1$.

For any $j \in \{1, 2, ..., m\}$, the sequence $\left\{\frac{\psi_{j;n}}{1-\psi_{j;n}} \times \frac{s_j}{1-s_j}\right\}_{n=1}^{\infty}$ is bounded and monotone decreasing, thus convergent. Let $\mathcal{L}_j$ be the limit of this sequence. If $\mathcal{L}_m < \frac{c}{1-c}$, by Lemma 1.13, there is some $N' \in \mathbb{N}$ such that for all $n \geq N'$, the most selective equilibrium is more selective than σ_m. So, buyers must be indifferent when they trade—expected trade surplus must be 0. Since total surplus is bounded below by $\Pi^{\emptyset}$, we conclude that $\Pi^n\left(\hat{\sigma}; \mathcal{E}\right) = \Pi^{\emptyset} = 0$ for all $n \geq N'$.

Now consider the case $\mathcal{L}_m \geq \frac{c}{1-c}$. Since $\psi_{j:n}$ is decreasing in n, the most selective equilibrium with n buyers, $\hat{\sigma}_n$, must get weakly more selective with n. So, the sequence $\{r_\theta(\hat{\sigma}_n; \mathcal{E})\}_{n=1}^{\infty}$ is weakly increasing in n, convergent, and Cauchy. If for any $N \in \mathbb{N}$, $\hat{\sigma}_N$ is more selective than σ_m, we are done by the argument in the preceding paragraph. Otherwise, the sequence $\{r_\theta(\hat{\sigma}_n; \mathcal{E})\}_{n=1}^{\infty}$ converges to a number below 1. If are both constant at 0, we are done—total surplus is equal to that under the no-information benchmark along the sequence. Otherwise, $r_L(\hat{\sigma}_n; \mathcal{E}) > r_H(\hat{\sigma}_n; \mathcal{E})$ along the sequence. Since both sequences

are Cauchy, there exists some $N \in \mathbb{N}$ and $M \geq N$ such that for all $m \geq M$, we have:

$$\begin{aligned}
&\rho \times (1-c) \times [1 - r_H(\hat{\sigma}_m; \mathcal{E})]^m - (1-\rho) \times c \times [1 - r_L(\hat{\sigma}_m; \mathcal{E})^m] \\
\approx &\rho \times (1-c) \times [1 - r_H(\hat{\sigma}_N; \mathcal{E})]^m - (1-\rho) \times c \times [1 - r_L(\hat{\sigma}_N; \mathcal{E})^m] \\
> &\rho \times (1-c) \times [1 - r_H(\hat{\sigma}_N; \mathcal{E})]^{m+1} - (1-\rho) \times c \times \left[1 - r_L(\hat{\sigma}_N; \mathcal{E})^{m+1}\right] \\
\approx &\rho \times (1-c) \times [1 - r_H(\hat{\sigma}_{m+1}; \mathcal{E})]^{m+1} - (1-\rho) \times c \times \left[1 - r_L(\hat{\sigma}_{m+1}; \mathcal{E})^{m+1}\right]
\end{aligned}$$

This proves that total surplus is eventually decreasing.

Furthermore, since there is at least one outcome of the experiment $\mathcal{E}$ where a buyer surely trades with the seller, as $n \to \infty$, the seller trades almost surely regardless of quality. Hence, total surplus converges to $\rho - c$. Since total surplus can never be negative, it must be that $\Pi^{\emptyset} = \rho - c$ in this case.

□

Theorem 1.2. Let buyers' experiment $\mathcal{E}$ be binary. Then, equilibrium* total surplus is increasing in the strength of good news (s_H) but is quasiconcave and eventually decreasing in the strength of bad news (s_L).

I will use Lemmas 1.15, 1.17, and 1.18 below, of independent interest, to prove Theorem 1.2. Throughout, I denote the most and least selective equilibrium strategies under the experiment $\mathcal{E}$ as $\hat{\sigma}^*_{\mathcal{E}}$ and $\check{\sigma}^*_{\mathcal{E}}$, respectively. I drop the subscript whenever the experiment in question is obvious.

Lemma 1.15. Let $\mathcal{E}$ be a binary experiment, with outcomes in $\mathbf{S} = \{s_L, s_H\}$; $s_L \leq s_H$. $\Psi(\sigma; \mathcal{E})$ is:

i strictly increasing in $\sigma(s_L)$, whenever $\sigma(s_H) = 1$,

ii strictly decreasing in $\sigma(s_H)$ whenever $\sigma(s_L) = 0$.

Proof. Part i:

Let $\sigma(s_L) \in (0,1)$ and $\sigma(s_H) = 1$. The interim belief $\Psi(\sigma;\mathcal{E})$ is then given by:

$$\begin{aligned}\Psi(\sigma;\mathcal{E}) &= \mathbb{P}\left(\theta = H \mid \text{visit received}\right) \\ &= \sum_{i=0}^{n-1} \mathbb{P}(\text{visited after i}^{\text{th}} \text{ rejection} \mid \text{visit received}) \times \mathbb{E}\left[\theta = H \mid \text{i } s_L \text{ signals}\right] \\ &= \sum_{i=0}^{n-1} \frac{\mathbb{P}(\text{visited after i}^{\text{th}} \text{ rejection})}{\mathbb{P}(\text{visit received})} \times \mathbb{E}\left[\theta = H \mid \text{i } s_L \text{ signals}\right]\end{aligned}$$

Note that $\mathbb{E}\left[\theta = H \mid \text{i } s_L \text{ signals}\right] < \mathbb{E}\left[\theta = H \mid \text{i+1 } s_L \text{ signals}\right]$; since every s_L signal is further evidence for $\theta = L$. We have:

$$\begin{aligned}\mathbb{P}(\text{ visited after i}^{\text{th}} \text{ rejection}) = &\mathbb{P}\left(\text{buyer was (i+1)}^{\text{st}} \text{ in order } \mid \text{seller got i rejections}\right) \\ &\times \mathbb{P}\left(\text{seller got i rejections}\right) \\ = &\frac{1}{n} \times \mathbb{P}(\text{i } s_L \text{ signals}) \times \left[1 - \sigma(s_L)\right]^i\end{aligned}$$

The proof is completed by noting that:

$$\frac{\mathbb{P}(\text{ visited after (i+1)}^{\text{st}} \text{ rejection})}{\mathbb{P}(\text{ visited after i}^{\text{th}} \text{ rejection})} = \frac{\mathbb{P}(\text{i+1 } s_L \text{ signals})}{\mathbb{P}(\text{i } s_L \text{ signals})} \times \left[1 - \sigma(s_L)\right]$$

decreases, and thus $\Psi(\sigma;\mathcal{E})$ increases, in $\sigma(s_L)$.

Part ii:

Now take $\sigma(s_L) = 0$. We then have:

$$r_H(\sigma;\mathcal{E}) = 1 - p_H(s_H)\sigma(s_H) \qquad\qquad r_L(\sigma;\mathcal{E}) = 1 - p_L(s_H)\sigma(s_H)$$

and:

$$\begin{aligned}\Psi(\sigma;\mathcal{E}) &\propto \frac{1 + r_H + ... + r_H^{n-1}}{1 + r_L + ... + r_L^{n-1}} \\ &= \frac{1 - r_H^n}{1 - r_L^n} \times \frac{1 - r_L}{1 - r_H} = \frac{1 - r_H^n}{1 - r_L^n} \times \frac{p_L(s_H)}{p_H(s_H)} \\ &\propto \frac{1 - r_H^n}{1 - r_L^n} = \frac{1 - (1 - p_H(s_H)\sigma(s_H))^n}{1 - (1 - p_L(s_H)\sigma(s_H))^n}\end{aligned}$$

Differentiating the last expression with respect to $\sigma(s_H)$ and rearranging its terms reveals

that this derivative is proportional to:

$$\frac{s_H}{1-s_H} \times \left(\frac{r_H}{r_L}\right)^{n-1} - \frac{1-(r_H)^n}{1-(r_L)^n}$$

The positive term is the likelihood ratio of one s_H signal and $n-1$ rejections, and the negative term is the likelihood ratio from *at most* $n-1$ rejections. Since acceptances only happen with s_H signals, the negative term strictly exceeds the positive term. This can be verified directly, too:

$$\begin{aligned}\frac{1-(r_H)^n}{1-(r_L)^n} > \frac{s_H}{1-s_H} \times \left(\frac{r_H}{r_L}\right)^{n-1} &\iff \frac{1-(r_H)^n}{1-(r_L)^n} \times \frac{1-r_L}{1-r_H} > \left(\frac{r_H}{r_L}\right)^{n-1} \\ &\iff \frac{1+...+(r_H)^{n-1}}{1+...+(r_L)^{n-1}} > \left(\frac{r_H}{r_L}\right)^{n-1}\end{aligned}$$

The last inequality can be verified easily. Thus, $\Psi(\sigma;\mathcal{E})$ decreases in $\sigma(s_H)$.

□

The Corollary below follows from Lemma 1.15. Let both $\mathcal{E}'$ and $\mathcal{E}$ are binary experiments, where the former is Blackwell more informative than the latter. If, under both experiments, every buyer accepts upon "good news" and rejects upon "bad news", the interim belief under $\mathcal{E}'$ is lower.

Corollary 1.16. Let $\mathcal{E}'$ and $\mathcal{E}$ be two binary experiments, where the former is Blackwell more informative than the latter. Let the strategies σ' and σ for these respective experiments be defined as:

$$\sigma'(s') := \begin{cases} 0 & s' = s'_L \\ 1 & s' = s'_H \end{cases} \qquad \sigma(s) := \begin{cases} 0 & s = s_L \\ 1 & s = s_H \end{cases}$$

Then, $\Psi(\sigma';\mathcal{E}') \leq \Psi(\sigma;\mathcal{E})$.

Proof. Establishing that this holds for a pair $(\mathcal{E}',\mathcal{E})$ for which either (i) $s'_H > s_H$ and $s_L = s'_L$, or (ii) $s'_H = s_H$ and $s_L > s'_L$ suffices. I will only prove the first case, the second is analogous. Below I show that the outcome induced by σ under experiment $\mathcal{E}$ can be replicated by some strategy $\tilde{\sigma}$ under experiment $\mathcal{E}'$, where $\tilde{\sigma}(s_L) > 0$ and $\tilde{\sigma}(s_H) = 1$. Then, the desired conclusion follows from Lemma 1.15.

Take the pair $(\sigma, \mathcal{E})$. The probabilities that a buyer accepts or rejects trade, conditional on θ, is given by:

$$\frac{\mathbb{P}_\sigma\left(\text{rejected} \mid \theta = H\right)}{\mathbb{P}_\sigma\left(\text{rejected} \mid \theta = L\right)} = \frac{s_L}{1 - s_L} \qquad \frac{\mathbb{P}_\sigma\left(\text{accepted} \mid \theta = H\right)}{\mathbb{P}_\sigma\left(\text{accepted} \mid \theta = L\right)} = \frac{s_H}{1 - s_H}$$

For the pair $(\tilde{\sigma}, \mathcal{E}')$ where $\tilde{\sigma}(s'_H) = 1$, we have:

$$\frac{\mathbb{P}_{\tilde{\sigma}}\left(\text{rejected} \mid \theta = H\right)}{\mathbb{P}_{\tilde{\sigma}}\left(\text{accepted} \mid \theta = L\right)} = \frac{s_L}{1 - s_L} \qquad \frac{\mathbb{P}_{\tilde{\sigma}}\left(\text{accepted} \mid \theta = H\right)}{\mathbb{P}_{\tilde{\sigma}}\left(\text{accepted} \mid \theta = L\right)} = \frac{p'_H(s_H) + \tilde{\sigma}(s_L)p'_H(s_L)}{p'_L(s_H) + \tilde{\sigma}(s_L)p'_L(s_L)}$$

where $\{p'_L, p'_H\}$ are the distributions for the experiment $\mathcal{E}'$. It is easy to verify that the expression on the right falls from $\frac{s'_H}{1-s'_H}$ to 1 monotonically and continuously as $\tilde{\sigma}(s_L)$ rises from 0 to 1. Thus, there is a unique interior value of $\tilde{\sigma}(s_L)$ that replicates the outcome of $(\sigma; \mathcal{E})$.

□

Lemma 1.17. Let $\mathcal{E}$ be a binary experiment, with outcomes in $\mathbf{S} = \{s_L, s_H\}$; $s_L \leq s_H$. Let $\sigma^\star$ be buyers' equilibrium strategy, either in the most or least selective equilibrium. Then, $\sigma^\star(s_L) \in \{0, 1\}$.

Proof. I start by proving this for the least selective equilibrium; i.e. $\check{\sigma}^*(s_L) \in \{0, 1\}$. For s_L^{mute} defined in Definition 1.10, observe that when $s_L \geq s_L^{\text{mute}}$, $\sigma(s_L) = \sigma(s_H) = 1$ is an equilibrium; so we must have $\check{\sigma}^*(s_L) = 1$. The strategy σ defined by $\sigma(s_H) = \sigma(s_L) = 1$ gives rise to the interim belief $\Psi(\sigma; \mathcal{E}) = \rho$, which in turn renders approving upon the outcome s_L optimal. In turn, if $s_L < s_L^{\text{mute}}$, we must have $\sigma^*(s_L) = 1$ for any equilibrium strategy; since the equilibrium interim belief always lies below the prior belief (Proposition 1.1).

Now consider the most selective equilibrium strategy; $\hat{\sigma}^*$. For contradiction, let $1 > \hat{\sigma}^*(s_L) > 0$ and $\hat{\sigma}^*(s_H) = 1$. Lemma 1.15 establishes that the interim belief falls as $\sigma(s_L)$ falls; which implies there must be another, more selective equilibrium strategy σ^* such that $\sigma^*(s_L) = 0$ and $\sigma^*(s_H) = 1$.

□

Lemma 1.17 establishes that when their experiment $\mathcal{E}$ is binary, buyers *never* mix upon seeing "bad news", $s = s_L$, neither in the most nor the least selective equilibrium. Following

up, Lemma 1.18 establishes that a more informative binary experiment pushes buyers to reject upon bad news in both equilibria.

Lemma 1.18. Let $\mathcal{E}$ be a binary experiment, with outcomes in $\mathbf{S} = \{s_L, s_H\}$; $s_L \leq s_H$. Buyers' acceptance probabilities upon "bad news", $s = s_L$, in the least and most selective equilibrium strategies are given by:

$$\check{\sigma}^*(s_L) = \begin{cases} 1 & s_L \geq s_L^{\text{mute}} \\ 0 & s_L < s_L^{\text{mute}} \end{cases} \qquad \hat{\sigma}^*(s_L) = \begin{cases} 1 & s_L < s_L^{\dagger}(s_H) \\ 0 & s_L \geq s_L^{\dagger}(s_H) \end{cases}$$

where $s_L^{\dagger}(.)$ is an increasing function of s_H, and $s_L^{\dagger}(s_H) \geq s_L^{\text{safe}}$.

Proof. Note that there exists an equilibrium where $\sigma(s_L) = 1$ if and only if:

$$\frac{\rho}{1-\rho} \times \frac{s_L}{1-s_L} \geq \frac{c}{1-c}$$

which, combined with Lemma 1.17, proves the part of the Lemma for the selective equilibrium.

Now, define the strategies σ_0 as σ_1 as:

$$\sigma_0(s) = \begin{cases} 0 & s = s_L \\ 0 & s = s_H \end{cases} \qquad \sigma_1(s) = \begin{cases} 0 & s = s_L \\ 1 & s = s_H \end{cases}$$

A necessary and sufficient condition for an equilibrium σ^* where $\sigma^*(s_L) = 0$ to exist is:

$$\frac{\Psi(\sigma_1; \mathcal{E})}{1 - \Psi(\sigma_1; \mathcal{E})} \times \frac{s_L}{1-s_L} \leq \frac{c}{1-c}$$

Sufficiency follows since either:

$$\frac{\Psi(\sigma_0; \mathcal{E})}{1 - \Psi(\sigma_0; \mathcal{E})} \times \frac{s_H}{1-s_H} \leq \frac{c}{1-c}$$

which implies σ_0 is an equilibrium, or there is an equilibrium strategy σ^* such that $\sigma^*(s_L) = 0$ and $\sigma^*(s_H) > 0$ by Lemma 1.15. The condition is necessary, since any strategy that is less selective than σ_1 induces a higher interim belief, by Lemma 1.15.

By Corollary 1.16, whenever this necessary and sufficient condition holds for an experiment $\mathcal{E}$, it also holds for a (Blackwell) more informative experiment $\mathcal{E}'$. Moreover, whenever the low signals are rejected in the least selective equilibrium, they must be in the most selective equilibrium. This concludes the proof.

□

Proof, Theorem 1.2: By Lemma 1.18, Blackwell improving buyers' experiment shifts both their least selective and most selective equilibrium strategies once from *always* accepting trade to rejecting upon the low signal. By Lemma 1.14, this shift in buyers' strategy increases efficiency—and therefore each buyer's expected surplus.

Let $\{\sigma_\alpha\}_{\alpha\in[0,1]}$ be the family of strategies where buyers reject upon the low signal:

$$\sigma_\alpha(s) = \begin{cases} 0 & s = s_L \\ 1 & s = s_H \end{cases}$$

By Lemma 1.15, the interim belief ψ_α that the strategy σ_α induces is strictly decreasing in α. Thus, at most one of these can be an equilibrium strategy for a given experiment. Furthermore, whenever buyers' expected surplus from σ_1 is weakly positive, this must be the equilibrium strategy; decreasing α can only make approving upon the high signal *more* profitable. Hence, whenever buyers reject upon the low signal in equilibrium, efficiency is given by: $\Pi(\sigma^*;\mathcal{E}) = \max\{0, \Pi(\sigma_1;\mathcal{E})\}$. The Theorem then follows from the Claim below:

Claim. $\max\{0, \Pi(\sigma_1;\mathcal{E})\}$ is:

i weakly increasing in s_H whenever there is some equilibrium strategy σ^* s.t. $\sigma^*(s_L) = 0$.

ii hump-shaped in s_L. As s_L falls, it:

- weakly increases when $s_L \geq s_L^{as}$,
- weakly decreases when $s_L \leq s_L^{as}$

where s_L^{as} is defined implicitly as:

$$\frac{\rho}{1-\rho} \times \left(\frac{s_L^{\text{as}}}{1-s_L^{\text{as}}}\right)^{n-1} \times \frac{s_H}{1-s_H} = \frac{c}{1-c}$$

Proof of the Claim.

Part i. Increasing the strength of good news; i.e. s_H.

Let $\mathcal{E}$ and $\mathcal{E}'$ be two binary experiments with outcome sets $\mathbf{S} = \{s_L, s_H\}$ and $\mathbf{S'} = \{s'_L, s'_H\}$. The experiment $\mathcal{E}'$ carries *marginally stronger good news* than experiment $\mathcal{E}$:

$$s'_L = s_L \qquad\qquad s'_H = s_H + \delta$$

for some small δ such that $1 - s_H \geq \delta > 0$. I show that $\Pi(\sigma'_1; \mathcal{E}') > \Pi(\sigma_1; \mathcal{E})$; where σ'_1 is defined analogously to σ_1 for experiment $\mathcal{E}'$.

Step 1. Replicating $\mathcal{E}'$ with a signal pair $(s, \hat{s})$.

Rather than observing the outcome of experiment $\mathcal{E}'$, say a buyer initially observes her original signal s, and then potentially an additional auxilliary signal $\hat{s}$. The first signal she receives, s, records the outcome of $\mathcal{E}$. If the low outcome s_L materialises, the buyer observes no more information. If, however, the high outcome s_H materialises, she then observes the additional auxiliary signal $\hat{s}$. This auxiliary signal records the outcome of *another* binary experiment, $\hat{\mathcal{E}}$. The outcome of $\hat{\mathcal{E}}$ is independent both from s and anything else any other buyer observes. Conditional on the asset's quality θ, the distribution over its outcomes is given by the pmf $p_\theta(.)$:

$$\hat{p}_H(\hat{s_H}) = 1 - \varepsilon \times \frac{s_L}{1 - s_L} \qquad\qquad \hat{p}_L(\hat{s_H}) = 1 - \varepsilon \times \frac{s_H}{1 - s_H}$$

The evolution of the buyer's beliefs when she observes this signal pair is determined by the two likelihood ratios:

$$\frac{\mathbb{P}((s, \hat{s}) = (s_H, \hat{s_H}) \mid \theta = H)}{\mathbb{P}((s, \hat{s}) = (s_H, \hat{s_H}) \mid \theta = L)} = \frac{s_H}{1 - s_H} \times \frac{1 - \varepsilon \times \dfrac{s_L}{1 - s_L}}{1 - \varepsilon \times \dfrac{s_H}{1 - s_H}} \tag{1.9.3}$$

$$\frac{\mathbb{P}((s, \hat{s}) = (s_H, \hat{s_L}) \mid \theta = H)}{\mathbb{P}((s, \hat{s}) = (s_H, \hat{s_L}) \mid \theta = L)} = \frac{s_L}{1 - s_L} \tag{1.9.4}$$

Note that the likelihood ratio 1.9.3 increases continuously with ε.

The information from observing the pair $(s, \hat{s})$ as such is equivalent to observing the

outcome of experiment $\mathcal{E}'$, when:

$$\frac{s_H}{1-s_H} \times \frac{1-\varepsilon \times \dfrac{s_L}{1-s_L}}{1-\varepsilon \times \dfrac{s_H}{1-s_H}} = \frac{s_H+\delta}{1-(s_H+\delta)} \tag{1.9.5}$$

for our chosen (δ, ε). I choose ε to satisfy this equality for our δ. As such, ε becomes a continuously increasing function of δ. Furthermore, note that by varying ε between 0 and $\frac{1-s_H}{s_H}$, we can replicate *any* experiment $\mathcal{E}'$ with $s'_L = s_L$ and $1 \geq s'_H \geq s_H$.

Step 2. $\pi(\sigma'_1; \mathcal{E}') \geq \pi(\sigma_1; \mathcal{E})$.

The buyer who observes the signal pair $(s, \hat{s})$ obtains equivalent information to that from $\mathcal{E}'$. We now must identify the strategy $\tilde{\sigma} : \{s_L, (s_H, \hat{s_H}), (s_H, \hat{s_L})\} \to [0,1]$ for this signal pair that replicates the outcome of the strategy σ'_1 for experiment $\mathcal{E}'$. This strategy is defined as:

$$\tilde{\sigma}(s_H, \hat{s_H}) = 1 \qquad\qquad \tilde{\sigma}(s_L) = \tilde{\sigma}(s_H, \hat{s_L}) = 0$$

and replicates the likelihood ratios of an acceptance and rejection signal under $\mathcal{E}'$.

Now, fix the seller's *signal profile* $\mathbf{s} = \{(s^i, \hat{s}^i)\}_{i=1}^n$ (defined in Section 1.9.1). I call a seller a *marginal admit* if his score profile is such that:

i for at least one $i \in \{1, 2, ..., n\}$, $s^i = s_H$, and

ii for *every* $i \in \{1, 2, ..., n\}$, either $s^i = s_L$, or $\hat{s}^i = \hat{s}_L$.

These marginal admits drive the wedge between efficiency under $\mathcal{E}'$ and $\mathcal{E}$: while some buyer trades under $\mathcal{E}$, they *all* reject him under $\hat{\mathcal{E}}$. So:

$$\Pi(\sigma'_1; \mathcal{E}') - \Pi(\sigma_1; \mathcal{E}) = \mathbb{P}\left(\text{marginal admit}\right) \times \underbrace{\left[c - \mathbb{P}\left(\theta = H \mid \text{marginal admit}\right)\right]}_{(1)}$$

A marginal admit only has signal realisations $(s, \hat{s}) = (s_H, \hat{s_L})$ or $s = s_L$. These carry equivalent information about θ. Thus, the expression (1) above equals:

$$c - \mathbb{P}\left[\theta = H \mid s^1 = ... = s^n = s_L\right]$$

In the relevant region where there is an equilibrium strategy that leads to rejections after the low outcome s_L, the expression above must be weakly positive:

$$\begin{aligned} c - \mathbb{P}\left[\theta = H \mid s^1 = ... = s^n = s_L\right] &\propto \frac{c}{1-c} \times \frac{\rho}{1-\rho} \times \left(\frac{s_L}{1-s_L}\right)^n \\ &\leq \frac{c}{1-c} - \frac{\rho}{1-\rho} \times \frac{\sum_{k=0}^{n-1} p_H(s_L)^k}{\sum_{k=0}^{n-1} p_L(s_L)^k} \times \frac{s_L}{1-s_L} \\ &= \frac{c}{1-c} - \frac{\Psi(\sigma_1;\mathcal{E})}{1-\Psi(\sigma_1;\mathcal{E})} \times \frac{s_L}{1-s_L} \leq 0 \end{aligned}$$

where the last inequality follows from the necessary and sufficient condition the proof of Lemma 1.18 introduced for such an equilibrium to exist.

Part ii. Increasing the strength of bad news; i.e. decreasing s_L.

Now, let the experiment $\mathcal{E}'$ carry *marginally stronger bad news* than experiment $\mathcal{E}$ instead; for some arbitrarily small $\delta \in [0, s_L]$:

$$s'_L = s_L - \delta \qquad\qquad s'_H = s_H$$

Where σ'_1 and σ_1 are defined as before, I show that:

i $\Pi(\sigma'_1;\mathcal{E}) - \Pi(\sigma_1;\mathcal{E}) \geq 0$ when $s_L \geq s_L^{as}$, and

ii $\Pi(\sigma'_1;\mathcal{E}) - \Pi(\sigma_1;\mathcal{E}) \leq 0$ when $s_L \leq s_L^{as}$

Step 1. Replicating $\mathcal{E}'$ with a signal pair $(s, \hat{s})$.

As before, let each buyer observe *two* signals, potentially: s and $\hat{s}$. She first observes s, which records the outcome of $\mathcal{E}$. If the high outcome s_H materialises, she receives no further information. If, however, the low outcome s_L materialises, she then observes the additional auxiliary signal $\hat{s}$, which records the outcome of *another* binary experiment, $\hat{\mathcal{E}}$. As before, the outcome of this experiment is independent both from s and anything observed by any other buyer. Its distribution conditional on the asset's quality θ is given by the pmf $p_\theta(.)$:

$$\hat{p}_H(\hat{s}_H) = \varepsilon \times \frac{s_H}{1-s_H} \qquad\qquad \hat{p}_L(\hat{s}_H) = \varepsilon \times \frac{s_L}{1-s_L}$$

The evolution of the buyer's beliefs upon seeing the signal pair $(s, \hat{s})$ is then determined by the two likelihood ratios:

$$\frac{\mathbb{P}\left((s,\hat{s}) = (s_L, \hat{s_H}) \mid \theta = H\right)}{\mathbb{P}\left((s,\hat{s}) = (s_L, \hat{s_H}) \mid \theta = H\right)} = \frac{s_H}{1-s_H} \tag{1.9.6}$$

$$\frac{\mathbb{P}\left((s,\hat{s}) = (s_L, \hat{s_L}) \mid \theta = H\right)}{\mathbb{P}\left((s,\hat{s}) = (s_L, \hat{s_L}) \mid \theta = H\right)} = \frac{s_L}{1-s_L} \times \frac{1-\varepsilon \times \dfrac{s_H}{1-s_H}}{1-\varepsilon \times \dfrac{s_L}{1-s_L}} \tag{1.9.7}$$

Note that 1.9.7 is continuously and strictly decreasing with ε, taking values between $\frac{s_L}{1-s_L}$ and 0 as ε varies between 0 and $\frac{s_H}{1-s_H}$. The signal pair $(s, \hat{s})$ is informationally equivalent to $\mathcal{E}'$ when:

$$\frac{s_L}{1-s_L} \times \frac{1-\varepsilon \times \dfrac{s_H}{1-s_H}}{1-\varepsilon \times \dfrac{s_L}{1-s_L}} = \frac{s_L - \delta}{1-(s_L-\delta)}$$

I choose ε to satisfy this equality. As before, ε then becomes a continuously increasing function of δ.

Step 2. $\Pi(\sigma_1'; \mathcal{E}') - \Pi(\sigma_1; \mathcal{E})$ $\begin{cases} \geq 0 & s_L \geq s_L^{\text{as}} \\ \leq 0 & s_L \leq s_L^{\text{as}} \end{cases}$

The buyer who observes the signal pair $(s, \hat{s})$ obtains equivalent information to that from $\mathcal{E}'$. We now must identify the strategy $\tilde{\sigma} : \{(s_L, \hat{s_H}), (s_L, \hat{s_L}), s_H\} \to [0,1]$ for this signal pair that replicates the outcome of the strategy σ_1' for experiment $\mathcal{E}'$. This strategy is defined as:

$$\tilde{\sigma}(s_L, \hat{s_H}) = \tilde{\sigma}(s_H) = 1 \qquad\qquad \tilde{\sigma}(s_L, \hat{s_L}) = 0$$

and replicates the likelihood ratios of an approval and rejection signal under $\mathcal{E}'$.

Now, fix the seller's *score profile*: $\mathbf{s} = \{(s^i, \hat{s}^i)\}_{i=1}^n$. I call a seller a *marginal reject* if:

i for every $i \in \{1, 2, ..., n\}$, $s^i = s_L$, and

ii for at least one $i \in \{1, 2, ..., n\}$, $\hat{s}^i = \hat{s}_H$.

Marginal rejects drive the wedge between efficiency under $\mathcal{E}'$ and $\mathcal{E}$: while *no* buyer trades

under $\mathcal{E}$, *at least one* buyer does under $\mathcal{E}'$. So:

$$\Pi(\sigma_1';\mathcal{E}') - \Pi(\sigma_1;\mathcal{E}) = \mathbb{P}\left(\text{marginal reject}\right) \times \underbrace{\left[\mathbb{P}\left(\theta = H \mid \text{marginal reject}\right) - c\right]}_{(2)}$$

For a marginal reject, buyers observe either $(s^i, \hat{s}^i) = (s_L, \hat{s}_L)$, or $(s^i, \hat{s}^i) = (s_L, \hat{s}_H)$. Denote the number of buyers who observed the latter as $\#$. Since the seller is a marginal reject, $\# \geq 1$. Then, (2) equals:

$$\underbrace{\sum_{i=1}^{n} \frac{\mathbb{P}\left(i \ \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)}{\sum_{j=1}^{n} \mathbb{P}\left(j \ \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)}}_{(3)} \times \mathbb{P}\left(\theta = H \mid \# = i\right) - c$$

where:

$$\mathbb{P}\left(i \ \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right) = k \times \binom{n}{i} \times \left(\frac{s_H}{1-s_H} \times \varepsilon\right)^i \times \left(1 - \frac{s_H}{1-s_H} \times \varepsilon\right)^{n-i}$$
$$+ (1-k) \times \binom{n}{i} \times \left(\frac{s_L}{1-s_L} \times \varepsilon\right)^i \times \left(1 - \frac{s_L}{1-s_L} \times \varepsilon\right)^{n-i}$$

and $k = \mathbb{P}\left(\theta = H \mid s^1 = ... = s^n = s_L\right)$. Thus, the limit of expression (3) as $\varepsilon \to 0$ (and therefore, $\delta \to 0$) for any $i > 1$ is:

$$\lim_{\varepsilon \to 0} \frac{\frac{1}{\varepsilon} \times \mathbb{P}\left(i \ \hat{s} = \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)}{\frac{1}{\varepsilon} \times \sum_{j=1}^{n} \mathbb{P}\left(j \ \hat{s} = \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)} = 0 \tag{1.9.8}$$

Therefore, we get:

$$\lim_{\varepsilon \to 0} \sum_{i=1}^{n} \frac{\mathbb{P}\left(i \ \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)}{\sum_{j=1}^{n} \mathbb{P}\left(j \ \hat{s_H} \text{ signals } \mid s^1 = ... = s^n = s_L\right)} \times \mathbb{P}\left(\theta = H \mid \# = i\right) - c$$
$$= \mathbb{P}\left(\theta = H \mid \# = 1\right) - c$$
$$\propto \frac{\rho}{1-\rho} \times \left(\frac{s_L}{1-s_L}\right)^{n-1} \times \frac{s_H}{1-s_H} - \frac{c}{1-c}$$

proving the claim. □

□

Proposition 1.4. Let buyers' experiment be binary. Then, equilibrium* total surplus weakly decreases with stronger bad news (lower s_L) when:

$$\frac{\rho}{1-\rho} \times \max\left\{\frac{s_L}{1-s_L}, \left[\frac{s_L}{1-s_L}\right]^{n-1} \times \frac{s_H}{1-s_H}\right\} \leq \frac{c}{1-c}$$

This condition is also necessary in the least selective equilibrium.

Proof.

i The least selective equilibrium:

By Lemma 1.17, the probability that the seller trades upon the low outcome in the least selective equilibrium is:

$$\check{\sigma}^*(s_L) = \begin{cases} 1 & s_L \geq s_L^{\text{mute}} \\ 0 & s_L < s_L^{\text{mute}} \end{cases}$$

Thus, efficiency equals (i) the expected surplus from always approving the applicant when $s_L \geq s_L^{\text{mute}}$, and (ii) $\max\{0, \Pi(\sigma_1; \mathcal{E})\}$ when $s_L < s_L^{\text{mute}}$ (established in the proof of Theorem 1.2):

$$\Pi(\check{\sigma}^*; \mathcal{E}) = \begin{cases} \rho - c & s_L \geq s_L^{\text{mute}} \\ \max\{0, \Pi(\sigma_1; \mathcal{E})\} & s_L < s_L^{\text{mute}} \end{cases}$$

Since always trading is always feasible, we have $\max\{0, \Pi(\sigma_1; \mathcal{E})\} \geq \rho - c$ when $s_L < s_L^{\text{mute}}$ by Lemma 1.14. Furthermore, the final Claim in Theorem 1.2's proof establishes that as s_L falls, the expression $\max\{0, \Pi(\sigma_1; \mathcal{E})\}$ weakly increases (decreases) when $s_L \geq s_L^{\text{as}}$ ($s_L \leq s_L^{\text{as}}$). Thus the desired conclusion is established.

ii The most selective equilibrium:

By Lemma 1.18, the most selective equilibrium shifts from one where a buyer always trades to one where she rejects upon the low signal when $s_H \geq s_H^\dagger(s_L)$, where $s_H^\dagger(.)$ is an increasing function of s_L. Following the arguments made for the least selective equilibrium then, efficiency:

- weakly increases as s_L decreases, when $s_L \geq \min\left\{s_L^{\text{as}}, s_L^\dagger(s_H)\right\}$

- weakly decreases as s_L decreases, when $s_L \leq \min\left\{s_L^{\text{as}}, s_L^{\dagger}(s_H)\right\}$.

The desired result follows by noting that $s_L^{\dagger}(s_H) \geq s_L^{\text{safe}}$, and therefore $\min\left\{s_L^{\dagger}, s_L^{\text{as}}(s_H)\right\} \geq \min\left\{s_L^{\text{safe}}, s_L^{\text{as}}(s_H)\right\}$.

□

Theorem 1.3. Let the experiment $\mathcal{E}'$ differ from $\mathcal{E}$ by a local spread at $s_j \in \mathbf{S}$. Equilibrium* total surplus is:

1. weakly greater under $\mathcal{E}'$ if the local spread is a negative override under equilibrium*.
2. weakly less under $\mathcal{E}'$ if the local spread is a positive override under equilibrium*, unless adverse selection is $\sigma^\star_{\mathcal{E}}$-irrelevant for signal s_{j+1}.

Proof. The Theorem focuses either on the least, or the most selective equilibrium strategies under both experiments. In the discussion below, I let σ^* and $\sigma^{*'}$ denote whichever equilibria we are focusing on under the respective experiments $\mathcal{E}$ and $\mathcal{E}'$. When I need to distinguish between the least and most selective equilibria, I denote them as $(\check{\sigma}, \check{\sigma}')$ and $(\hat{\sigma}, \hat{\sigma}')$, respectively. Following the notation introduced in Definition 1.2, let $\mathbf{S} \cup \mathbf{S}' = \{s_1, s_2, ..., s_M\}$ be the joint support of the experiments $\mathcal{E}$ and $\mathcal{E}'$, with elements increasing in their indices as usual. Since $\mathcal{E}'$ is obtained by a *local* mean preserving spread of $\mathcal{E}$, there is a monotone strategy $\sigma' : \mathbf{S}' \to [0, 1]$ whose outcome under $\mathcal{E}'$ replicates the outcome of σ^* under $\mathcal{E}$:

$$\sigma'(s) = \begin{cases} \sigma^*(s_j) & s \in \{s_{j-1}, s_{j+1}\} \\ \sigma^*(s) & s \notin \{s_{j-1}, s_{j+1}\} \end{cases}$$

Claim 1. Efficiency under the most (least) selective equilibrium of $\mathcal{E}'$ weakly exceeds that under $\mathcal{E}$ when $\hat{\sigma}(s_j) = 1$ $(\check{\sigma}(s_j) = 1)$.

Now suppose s_j leads to trade under σ^*; $\sigma^*(s_j) = 1$. Therefore, $\sigma'(s_{j-1}) = \sigma'(s_{j+1}) = 1$. Below, I show that $\sigma^{*'}$ is *more selective than* σ'. By Lemma 1.11, it follows that $\Pi\left(\sigma^{*'}; \mathcal{E}'\right) \geq \Pi\left(\sigma'; \mathcal{E}'\right) = \Pi\left(\sigma; \mathcal{E}\right)$.

If $s_{j-1} = \min \mathbf{S} \cup \mathbf{S}'$ or $\sigma^{*'}(s_{j-2}) = 0$, $\sigma^{*'}$ must necessarily be more selective than σ'; and we are done. So, for contradiction, I assume the following:

- $s_{j-1} > \min \mathbf{S} \cup \mathbf{S}'$

- $\sigma^{*'}(s_{j-2}) > 0$
- $\sigma^{*'}$ is *less* selective than σ', where the two strategies differ.

Case i. σ^* and $\sigma^{*'}$ are the least selective equilibrium strategies; i.e. $\sigma^* = \check{\sigma}$ and $\sigma^{*'} = \check{\sigma}'$.

I will prove the contradiction by constructing a strategy $\tilde{\sigma} : \mathbf{S} \to [0, 1]$ for experiment $\mathcal{E}$, such that:

i $\tilde{\sigma}$ replicates the outcome $\check{\sigma}'$ induces in $\mathcal{E}'$,

ii That $\check{\sigma}'$ is an eqm. strategy under $\mathcal{E}'$ implies that $\tilde{\sigma}$ is an eqm. strategy under $\mathcal{E}$,

iii But $\tilde{\sigma}$ is less selective than $\check{\sigma}$, contradicting that $\check{\sigma}$ is the least selective equilibrium strategy under $\mathcal{E}$.

I define the strategy $\tilde{\sigma} : \mathbf{S} \to [0, 1]$ for $\mathcal{E}$ as:

$$\tilde{\sigma}(s) := \begin{cases} 1 & s = s_i \\ \sigma'(s) & s \neq s_i \end{cases}$$

it is seen easily that $\tilde{\sigma}$ replicates the outcome of $\check{\sigma}'$. Furthermore, $\check{\sigma}'$ is an equilibrium under $\mathcal{E}'$ if and only if $\tilde{\sigma}$ is an equilibrium under $\mathcal{E}$: they induce the same interim belief ψ, and share the following necessary and sufficient condition for optimality:

$$\mathbb{P}_\psi\left(\theta = H \mid s_{j-2}\right) \begin{cases} = c & \sigma'(s_{j-2}) < 1 \\ \geq c & \sigma'(s_{j-2}) = 1 \end{cases}$$

The strategy $\tilde{\sigma}$ under experiment $\mathcal{E}$ replicates the outcome of $\check{\sigma}'$ under experiment $\mathcal{E}'$, and σ' under $\mathcal{E}'$ replicates the outcome of $\check{\sigma}$ under experiment $\mathcal{E}$. Since we assumed that $\check{\sigma}'$ is less selective than σ', it must be that $\tilde{\sigma}$ is less selective than $\check{\sigma}$.

Case ii. σ^* and $\sigma^{*'}$ are the most selective equilibrium strategies; i.e. $\sigma^* = \hat{\sigma}$ and $\sigma^{*'} = \hat{\sigma}'$.

Since strategy σ' for experiment $\mathcal{E}'$ replicates the outcome of $\hat{\sigma}$ for experiment $\mathcal{E}$, the two strategies induce the same interim belief ψ. Therefore, if $\mathbb{P}_\psi\left(\theta = H \mid s_{j-1}\right) \geq c$, σ' is an equilibrium under $\mathcal{E}'$; meaning $\hat{\sigma}'$ must be more selective than σ'.

Otherwise, say $\mathbb{P}_\psi(\theta = H \mid s_{j-1}) < c$. Then, by Lemma 1.13, there must be an equilibrium strategy that is more selective than σ' under $\mathcal{E}'$.

Claim 2. Efficiency under the most (least) selective equilibrium of $\mathcal{E}'$ falls weakly below that under $\mathcal{E}$ if:

i. s_j leads to rejections under $\mathcal{E}$; i.e. $\hat{\sigma}(s_j) = 0$ ($\check{\sigma}(s_j) = 0$), and

ii. the following condition holds:

$$\frac{\rho}{1-\rho} \times \left(\frac{r_H(\sigma;\mathcal{E})}{r_L(\sigma;\mathcal{E})}\right)^{n-1} \times \frac{s_{j+1}}{1-s_{j+1}} \leq \frac{c}{1-c}$$

Now, suppose s_j leads to rejections under σ^*; $\sigma^*(s_j) = 0$. Consequently, we have $\sigma'(s_{j-1}) = \sigma'(s_{j+1}) = 0$. I establish Claim 2 in two steps:

Step 1: $\sigma^{*'}$ is less selective than σ'; trade is likelier when when s_j is locally spread.

Step 2: This efficiency when the condition in Claim 2 is met; $\Pi(\sigma^{*'};\mathcal{E}') \leq \Pi(\sigma';\mathcal{E}') = \Pi(\sigma^*;\mathcal{E})$.

Step 1.

If $s_{j+1} = \max \mathbf{S} \cup \mathbf{S}'$ or $\sigma^{*'}(s_{j+1}) > 0$, it must be the case that $\sigma^{*'}$ is less selective than σ', and we are done. So instead, I assume that $s_{j+1} < \max \mathbf{S} \cup \mathbf{S}'$ and $\sigma^{*'}(s_{j+1}) = 0$.

Case i. σ^* and $\sigma^{*'}$ are the least selective equilibrium strategies; i.e. $\sigma^* = \check{\sigma}$ and $\sigma^{*'} = \check{\sigma}'$

Since σ' replicates the outcome of $\check{\sigma}$, we have $\Psi(\check{\sigma};\mathcal{E}) = \Psi(\sigma';\mathcal{E}') = \psi$. Thus, σ' must be an equilibrium strategy under $\mathcal{E}'$ if $\mathbb{P}_\psi(\theta = H \mid s_{j+1}) \leq c$: the optimality conditions for all signals below s_{j+1} are satisfied *a fortiori*, and those for the signals above s_{j+1} are satisfied since $\check{\sigma}$ has the same optimality conditions under $\mathcal{E}$. So, $\check{\sigma}'$ must be less selective than σ', since the former is the least selective equilibrium. If on the other hand, $\mathbb{P}_\psi(\theta = H \mid s_{j+1}) > c$, there must be an equilibrium strategy under experiment $\mathcal{E}'$ that is *less* selective than σ', by Lemma 1.13.

Case ii. σ^* and $\sigma^{*'}$ are the most selective equilibrium strategies; i.e. $\sigma^* = \hat{\sigma}$ and $\sigma^{*'} = \hat{\sigma}'$.

$\hat{\sigma}'$ is the most selective equilibrium strategy under experiment $\mathcal{E}'$, and we assumed that $\hat{\sigma}'(s_{j+1}) = 0$. The strategy $\tilde{\sigma}$ defined below for experiment $\mathcal{E}$ replicates the outcome $\hat{\sigma}'$

generates under experiment $\mathcal{E}'$:

$$\tilde{\sigma}(s) = \begin{cases} 0 & s \leq s_j \\ \hat{\sigma}'(s) & s > s_j \end{cases}$$

Note that $\tilde{\sigma}$ must be an equilibrium under experiment $\mathcal{E}$, since the interim belief it induces is the same as the one $\hat{\sigma}'$ does, and its optimality constraints are a subset of the latter's. But since $\hat{\sigma}$ is the *most* selective equilibrium strategy under $\mathcal{E}$, $\tilde{\sigma}$ must be less selective than it.

Step 2.

The statement is trivially true when $\sigma' = \sigma^{*'}$, so I focus on the case where these two strategies differ. As Step 1 established, $\sigma^{*'}$ must be less selective than σ'. This implies that $\sigma^{*'}(s_{j+1}) > 0$. To see why, say we had $\sigma^{*'}(s_{j+1}) = 0$ instead. We can then construct a strategy $\tilde{\sigma}$ for experiment $\mathcal{E}$, which replicates the outcome $\sigma^{*'}$ generates under experiment $\mathcal{E}'$:

$$\tilde{\sigma}(s) = \begin{cases} 0 & s \leq s_j \\ \sigma^{*'}(s) & s > s_j \end{cases}$$

As they induce the same interim belief and the optimality constraints of the latter are a subset of the former's, $\tilde{\sigma}$ must be an equilibrium under $\mathcal{E}$. This contradicts with σ^* and $\sigma^{*'}$ being the least selective strategies; since $\sigma^{*'}$ being less selective than σ' implies that $\tilde{\sigma}$ must be less selective than σ^*. It also contradicts with σ^* and $\sigma^{*'}$ being the most selective strategies; since it would imply that σ', more selective than $\sigma^{*'}$, should be an equilibrium under $\mathcal{E}'$.

Given that $\sigma^{*'}(s_{j+1}) > 0$, I now take another strategy $\sigma^{\delta}_{\mathcal{E}'} : \mathbf{S}' \to [0,1]$ for experiment $\mathcal{E}'$:

$$\sigma^{\delta}_{\mathcal{E}'}(s) = \begin{cases} 1 & s > s_{j+1} \\ \delta & s = s_{j+1} \\ 0 & s < s_{j+1} \end{cases}$$

where $\delta > 0$ is small enough so that $\sigma^{\delta}_{\mathcal{E}'}$ is more selective than $\sigma^{*'}$, but less selective than σ'.

I will show that, when the condition stated in Claim 2 holds, we have $\Pi\left(\sigma^{\delta}_{\mathcal{E}'};\mathcal{E}'\right) \leq \Pi\left(\sigma';\mathcal{E}'\right)$ for $\delta \to 0$. Lemma 1.14 then implies that $\Pi\left(\sigma^{*'};\mathcal{E}'\right) \leq \Pi\left(\sigma^{\delta}_{\mathcal{E}'};\mathcal{E}'\right)$, which coins the result.

To show this, I construct another experiment $\mathcal{E}^{\text{re}}$ under which I will use compare two strategies, σ_{re} and $\sigma^{\delta}_{\text{re}}$, that replicate the outcomes of the strategies σ' and $\sigma^{\delta}_{\mathcal{E}'}$, respectively. The experiment $\mathcal{E}^{\text{re}}$ has three possible outcomes, $\{s^{\text{re}}_L, s^{\text{re}}_{\delta}, s^{\text{re}}_H\}$. Conditional on the applicant's quality θ, its outcome distribution is independent from any other information any evaluator sees, and is given by the following pmf p^{re}_{θ}:

$$p_{\theta}(s^{\text{re}}) = \begin{cases} 1 - r_{\theta}(\sigma^*;\mathcal{E}) & s = s^{\text{re}}_H \\ \delta \times p'_{\theta}(s_{j+1}) & s = s^{\text{re}}_{\delta} \\ r_{\theta}\left(\sigma^*;\mathcal{E}\right) - \delta \times p'_{\theta}(s_{j+1}) & s = s^{\text{re}}_L \end{cases}$$

Define the strategies σ_{re} and $\sigma^{\delta}_{\text{re}}$ for this experiment as follows:

$$\sigma_{\text{re}}(s) = \begin{cases} 1 & s = s^{\text{re}}_H \\ 0 & s = s^{\text{re}}_{\delta} \\ 0 & s = s^{\text{re}}_L \end{cases} \qquad \sigma^{\delta}_{\text{re}} = \begin{cases} 1 & s = s^{\text{re}}_H \\ 1 & s = s^{\text{re}}_{\delta} \\ 0 & s = s^{\text{re}}_L \end{cases}$$

Now note that these two strategies replicate the outcomes of the strategies σ' and $\sigma^{\delta}_{\mathcal{E}'}$, respectively. Under $\sigma_{\text{re}}(s)$, the probability that a buyer trades, conditional on the seller's quality, is the same as it is under strategy σ' (or σ^*, which it replicates), and under $\sigma^{\delta}_{\text{re}}$, it is the same as it is under $\sigma^{\delta}_{\mathcal{E}'}$.

So, the difference between efficiency under these two strategies is determined by the *marginal reject* who:

- is rejected by *every* buyer under the strategy σ_{re}.
- is accepted by *at least one* buyer under the strategy $\sigma^{\delta}_{\text{re}}$.

Where $\mathbf{s^{re}} = \left\{s^1, ..., s^n\right\}$ is the seller's signal profile under the experiment $\mathcal{E}^{\text{re}}$, he has:

- *no* s^{re}_H signals; $s^i \neq s^{\text{re}}_H$ for all $i \in \{1, 2, ..., n\}$ and
- *at least one* s^{re}_{δ} signal; there exists some $i \in \{1, 2, ..., n\}$ such that $s^i = s^{\text{re}}_H$.

Thus we have:

$$\Pi(\sigma^{\delta}_{\mathcal{E}'};\mathcal{E}') - \Pi(\sigma';\mathcal{E}') = \Pi(\sigma^{\delta}_{\text{re}};\mathcal{E}^{\text{re}}) - \Pi(\sigma_{\text{re}};\mathcal{E}^{\text{re}})$$
$$= \mathbb{P}\left(\text{marginal reject}\right) \times \underbrace{\left[\mathbb{P}\left(\theta = H \mid \text{marginal reject}\right) - c\right]}_{(2)}$$

The expression labelled (2) above equals:

$$\sum_{i=1}^{n} \frac{\mathbb{P}\left(i \ s^{\text{re}}_{\delta} \text{ and } n-i \ s^{\text{re}}_{L} \text{ signals}\right)}{\sum\limits_{k=1}^{n} \mathbb{P}\left(k \ s^{\text{re}}_{\delta} \text{ and } n-k \ s^{\text{re}}_{L} \text{ signals}\right)} \times \mathbb{P}\left(\theta = H \mid i \ s^{\text{re}}_{\delta} \text{ and } n-i \ s^{\text{re}}_{L} \text{ signals}\right) - c$$

Since the probability that a buyer observes the s^{re}_{δ} signal is proportional to δ, we have:[45]

$$\lim_{\delta \to 0} \frac{\mathbb{P}\left(i \ s^{\text{re}}_{\delta} \text{ and } n-i \ s^{\text{re}}_{L} \text{ signals}\right)}{\sum\limits_{k=1}^{n} \mathbb{P}\left(k \ s^{\text{re}}_{\delta} \text{ and } n-k \ s^{\text{re}}_{L} \text{ signals}\right)} = 0$$

Therefore, we get:

$$\lim_{\delta \to 0} \sum_{i=1}^{n} \frac{\mathbb{P}\left(i \ s^{\text{re}}_{\delta} \text{ and } n-i \ s^{\text{re}}_{L} \text{ signals}\right)}{\sum\limits_{k=1}^{n} \mathbb{P}\left(k \ s^{\text{re}}_{\delta} \text{ and } n-k \ s^{\text{re}}_{L} \text{ signals}\right)} \times \mathbb{P}\left(\theta = H \mid i \ s^{\text{re}}_{\delta} \text{ and } n-i \ s^{\text{re}}_{L} \text{ signals}\right) - c$$

$$\lim_{\delta \to 0} \mathbb{P}\left(\theta = H \mid \text{one } s^{\text{re}}_{\delta} \text{ signal and } n-1 \ s^{\text{re}}_{L} \text{ signals}\right) - c$$

$$\propto \lim_{\delta \to 0} \frac{\rho}{1-\rho} \times \frac{p'_H(s_{j+1})}{p'_L(s_{j+1})} \times \left(\frac{r_H(\sigma^*;\mathcal{E}) - \delta \times p'_H(s_{j+1})}{r_L(\sigma^*;\mathcal{E}) - \delta \times p'_L(s_{j+1})}\right)^{n-1} - \frac{c}{1-c}$$

$$= \frac{\rho}{1-\rho} \times \frac{s_{j+1}}{1-s_{j+1}} \left(\frac{r_H(\sigma^*;\mathcal{E})}{r_L(\sigma^*;\mathcal{E})}\right)^{n-1} - \frac{c}{1-c}$$

□

Proposition 1.6. Let the experiment $\mathcal{E}'$ differ from $\mathcal{E}$ by a local spread at s_j. Total surplus in the most selective equilibrium is lower under $\mathcal{E}'$ if the following conditions hold:

$$\frac{\rho}{1-\rho} \times \left(\frac{s_j}{1-s_j}\right) \leq \frac{c}{1-c} \qquad \text{and} \qquad \frac{\rho}{1-\rho} \times \left(\frac{s_j}{1-s_j}\right)^{n-1} \times \frac{s_{j+1}}{1-s_{j+1}} \leq \frac{c}{1-c}$$

Proof. First, I let $\hat{\sigma}(s_{j+2}) < 1$. I show that this implies $\hat{\sigma}$ and $\hat{\sigma}'$ induce equivalent

[45] See expression 1.9.8 and the surrounding discussion in the proof of Theorem 1.2 for a more detailed explanation of this.

outcomes under their respective experiments. The strategy $\sigma' : \mathbf{S}' \to [0,1]$ which replicates the outcome of $\hat{\sigma}'$ under experiment $\mathcal{E}'$:

$$\sigma'(s) = \begin{cases} \hat{\sigma}(s) & s \geq s_{j+2} \\ 0 & s < s_{j+2} \end{cases}$$

must then be an equilibrium strategy under experiment $\mathcal{E}'$. This is because these strategies induce the same interim belief, that $\hat{\sigma}$ is an equilibrium strategy under $\mathcal{E}$ ensures that the optimality conditions of σ' for signals below s_{j+2} are satisfied, and for signals above s_{j+2}, the optimality conditions are the same as those for σ'. This means that $\hat{\sigma}'$ must be more selective than σ'. However, when proving Theorem 1.3, we established that σ' must be more selective than $\hat{\sigma}'$. So it must be that $\sigma' = \hat{\sigma}'$, and we are done.

So instead, let $\hat{\sigma}(s_{j+2}) = 1$. But then, it is easily established that:

$$\frac{r_H(\hat{\sigma}; \mathcal{E})}{r_L(\hat{\sigma}; \mathcal{E})} \leq \frac{s_j}{1 - s_j}$$

since $r_\theta(\hat{\sigma}; \mathcal{E}) = \sum_{k=1}^{j} p_\theta(s_k)$. So, the condition Proposition 1.6 supplies is sufficient for the one Theorem 1.3 does.

□

Lemma 1.7. Where it exists, the regulator-optimal garbling is monotone binary and produces IC recommendations.

Proof. To prove this statement, I take some garbling $\mathcal{E}^G$ and an equilibrium $\sigma^G : \mathbf{S}^G \to [0,1]$ it supports. I then construct a monotone binary garbling $\mathcal{E}^{G*}$ which produces IC recommendations, and show that efficiency under $\mathcal{E}^{G*}$ and the strategy which obeys its recommendations, σ^{G*}, are higher than those under $\mathcal{E}^G$ and σ^G.

For the monotone binary garbling $\mathcal{E}^{G*} = (\mathbf{S}^G, \mathbf{P}^{G*})$ and the garbling $\mathcal{E}^G = (\mathbf{S}^G, \mathbf{P}^G)$ in question:

$$\mathbf{P} \times \mathbf{T} = \mathbf{P}^G \qquad\qquad \mathbf{P} \times \mathbf{T}^* = \mathbf{P}^{G*}$$

define the expressions:

$$f^*(s) := p_L(s) \times t^*_{i1} \qquad\qquad f(s) := p_L(s) \times \sum_{s^G_j \in \mathbf{S}^G} t_{ij} \times \left(1 - \sigma^G(s^G_j)\right)$$

for each $s \in \mathbf{S}$. Given the asset has Low quality, $f^*(s)$ is the probability that (i) a buyer would have observed the signal $s \in \mathbf{S}$ in her original experiment, *and* (ii) the garbling $\mathcal{E}^{G*}$ issues her a "rejection recommendation". Similarly, $f(s)$ is the probability that (i) a buyer would observe the signal $s \in \mathbf{S}$ in her original experiment, *and* (ii) he would be rejected under the equilibrium strategies σ^G. For this Low quality seller, r_L^{G*} below is the probability that the buyer receives a rejection recommendation under $\mathcal{E}^{G*}$; and r_L^G is the probability that the buyer rejects him under $(\mathcal{E}^G, \sigma^G)$:

$$r_L^{G*} := \sum_{s \in \mathbf{S}} f^*(s) \qquad\qquad r_L^G := \sum_{s \in \mathbf{S}} f(s)$$

Now, take the least selective monotone binary garbling $\mathcal{E}^{G*}$ such that $r_L^{G*} = r_L^G$. Evidently, this garbling exists.

Clearly, one can treat f^* and f as probability density functions over $\mathbf{S}$ when normalised. Furthermore, the distribution the former describes is first order stochastically dominated by the one the latter does; $\frac{f^*(s_j)}{\sum_{s \in \mathbf{S}} f^*(s)}$ crosses $\frac{f(s)}{\sum_{s \in \mathbf{S}} f(s)}$ once from below. Therefore we get:

$$\begin{aligned} r^*_H := \sum_{s \in \mathbf{S}} \frac{p_H(s)}{p_L(s)} \times \frac{f^*(s)}{\sum_{s \in \mathbf{S}} f^*(s)} & \\ \leq \sum_{s \in \mathbf{S}} \frac{p_H(s)}{p_L(s)} \times \frac{f(s)}{\sum_{s \in \mathbf{S}} f(s)} &=: r_H \end{aligned}$$

where r^*_H and r_H are the probabilities that a High quality seller is rejected from a visit under the strategies σ^{G*} and σ^G, respectively.

Since $r^*_H \geq r_H$ and $r^*_L = r_L$, efficiency is higher under σ^* than it is under σ. It only remains to show that the strategy σ^* is optimal against the interim belief ψ^* consistent with it.

The interim belief ψ^* consistent with $\mathcal{E}^{G*}$ and σ^{G*} lies below ψ—the interim belief

consistent with $\mathcal{E}^G$ and σ^G:

$$\frac{\psi^*}{1-\psi^*} = \frac{\sum_{k=0}^{n-1}(r_H^*)^k}{\sum_{k=0}^{n-1}(r_L^*)^k} = \frac{\sum_{k=0}^{n-1}(r_H)^k}{\sum_{k=0}^{n-1}(r_L)^k} \leq \frac{\psi}{1-\psi}$$

Under the interim belief ψ^*, it is optimal for a buyer upon the signal s_L^{G*} if and only if:

$$\frac{\psi^*}{1-\psi^*} \times \frac{r_H^*}{r_L^*} \leq \frac{c}{1-c}$$

But this inequality must hold; since $\frac{r_H^*}{r_L^*} \leq \frac{r_H}{r_L}$, $\psi^* \leq \psi$, and σ^G is optimal against ψ:

$$\frac{\psi^*}{1-\psi^*} \times \frac{r_H^*}{r_L^*} \leq \frac{\psi}{1-\psi} \times \frac{r_H}{r_L} \leq \frac{c}{1-c}$$

Furthermore, that $\Pi(\sigma^{G*}; S^{G*}) \geq \Pi\left(\sigma; S^G\right) \geq 0$ suggests that the expected suprlus from accepting a seller upon the "approve" recommendation must be weakly positive; hence optimal. Thus, the strategy σ^{G*} is optimal against ψ^*.

□

Proposition 1.8. If the least selective monotone binary garbling under which adverse selection is irrelevant produces IC recommendations, it is the regulator-optimal garbling. Otherwise, the regulator-optimal garbling is either:

- the least selective garbling under which adverse selection is irrelevant, or
- the most selective garbling under which adverse selection is not irrelevant

among monotone binary garblings which produce IC recommendations.

Proof. **Step 1:** The following are well-defined:

- the least selective monotone binary garbling under which adverse selection is irrelevant,
- the least (most) selective monotone binary garbling under which adverse selection is (not) irrelevant among those which produce IC recommendations.

I first show the least selective monotone binary garbling under which adverse selection is irrelevant is well defined. For any monotone binary garbling $\mathcal{E}^G = (\mathbf{S}^G, \mathbf{P}^G)$, let $\mathbf{P} \times \mathbf{T} =$

$\mathbf{P}^G$ and define $d(\mathcal{E}^G) := \sum_{i=1}^{m} t_{i2}$. Evidently, $d(.)$ is a bijection between the space of monotone binary garblings of $\mathcal{E}$ and $[0, m]$. Also, where both are monotone binary garblings of $\mathcal{E}$, $\mathcal{E}^G$ is more selective than $\mathcal{E}^{G'}$ if and only if $d(\mathcal{E}^G) \leq d(\mathcal{E}^{G'})$. Thus, we seek the monotone binary garbling $d^{-1}(D^*)$ where $D^* := \max\{D \in [0, m] : \text{ a.s. is irrelevant under } d^{-1}(D)\}$. We must only show that D^* is well defined. To that end, define the Real valued function F over the space of monotone binary garblings, where:

$$F(S^G) = \begin{cases} \frac{\rho}{1-\rho}\frac{p_H(s^*)}{p_L(s^*)} \times \left(\frac{r_H^G}{r_L^G}\right)^{n-1} - \frac{c}{1-c} & d(\mathcal{E}^G) \in (0, m) \\ \lim_{D \downarrow 0} F \circ d^{-1}(D) & d(\mathcal{E}^G) = 0 \qquad r_\theta^G := \sum_{s \in \mathbf{S}} p_s^G(s_L^G) \times p_\theta(s) \\ +\infty & d(\mathcal{E}^G) = m \end{cases}$$

where s^* is the threshold signal of this garbling.

So, we seek $D^* := \max\{D \in [0, m] : F \circ d^{-1}(D) \geq 0\}$. But this maximiser exists because the function $F \circ d^{-1}$ is upper semicontinuous: $F \circ d^{-1}$ is a decreasing function, and for any $\bar{D} \in [0, m]$ and $\varepsilon > 0$, we can find some δ_ε such that $D \in (\bar{D} - \delta_\varepsilon, \bar{D}) \cap [0, m]$ implies $d^{-1}(\bar{D})$ and $d^{-1}(D)$ have the same threshold signal s^* and thus $F \circ d^{-1}(D) < F \circ d^{-1}(\bar{D}) + \varepsilon$ since r_H^G/r_L^G is continuous in $d(S^G)$.

Now say this garbling does not provide IC recommendations. Denote the interim belief that is consistent with evaluators following $\mathcal{E}^G$'s recommendations as ψ^G. The garbling $\mathcal{E}^G$ provides IC recommendations if:

$$\underbrace{\frac{r_H^G}{r_L^G} \times \frac{\psi^G}{1-\psi^G}}_{:= f_1(d(\mathcal{E}^G))} \leq \frac{c}{1-c} \qquad \underbrace{\frac{1-r_H^G}{1-r_L^G} \times \frac{\psi^G}{1-\psi^G}}_{:= f_2(d(\mathcal{E}^G))} \geq \frac{c}{1-c}$$

As defined above, both $f_1(.)$ and $f_2(.)$ are continuous. Therefore, the set of monotone binary garblings with optimal recommendations—$\left\{D \in [0, m] : f_1(D) \leq \frac{c}{1-c} \text{ and } f_2(D) \geq \frac{c}{1-c}\right\}$—is compact. Thus, both objects below are well-defined:

$$\max\{D \in [0, m] \text{ and } d^{-1}(D) \text{ has IC rec.s } : F \circ d^{-1}(D) \geq 0\}$$
$$\min\{D \in [0, m] \text{ and } d^{-1}(D) \text{ has IC rec.s } : F \circ d^{-1}(D) \leq 0\}$$

Among those with IC recommendations, the former gives us the least selective garbling under which adverse selection is irrelevant. The latter gives us the most selective garbling under which adverse selection is not irrelevant among such garblings, since the least-selective garbling under which adverse selection is irrelevant does *not* have IC recommendations (the minimiser of this set *must* have $F \circ d^{-1}(D) < 0$).

Step 2: Proving Proposition 1.8 as stated.

Efficiency under a monotone binary garbling S^G and strategies σ^G that obey its recommendations is given by:

$$\Pi\left(\sigma^G; S^G\right) = \rho - c - \rho \times \left(r_H^G\right)^n \times (1-c) + (1-\rho) \times \left(r_L^G\right)^n \times c$$

As a function of $d^{-1}(.)$, efficiency is continuous and therefore attains its maximum over the set $[0, m]$. I show that this maximum is attained with the least selective garbling under which adverse selection is irrelevant.

For the garbling $\mathcal{E}^G$, define $\mathcal{E}^G_{+\varepsilon} := d^{-1}(d(\mathcal{E}^G)+\varepsilon)$ and $S^G_{-\varepsilon} := d^{-1}(d(S^G)-\varepsilon)$. Likewise, let $s^*_{+\delta}$ and $s^*_{-\delta}$ be the threshold signals of these experiments, and $r^*_{\theta;+\delta}$, $r^*_{\theta;-\delta}$ be the probability that a seller of quality θ is rejected in a visit, under each garbling. From our earlier reasoning about the impact of making evaluators strategies marginally more (less) selective, we observe that:

$$\lim_{\delta\to 0} \Pi\left(\sigma^G_{+\delta}; \mathcal{E}^G_{+\delta}\right) - \Pi\left(\sigma^G; \mathcal{E}^G\right) \propto \lim_{\delta\to 0} \frac{\rho}{1-\rho} \times \frac{p_H(s^*_{+\delta})}{p_L(s^*_{+\delta})} \times \left(\frac{r^G_{H;+\delta}}{r^G_{L;+\delta}}\right)^{n-1} - \frac{c}{1-c} \leq 0$$

where the last inequality follows since $\mathcal{E}^G$ is the least selective garbling under which adverse selection is irrelevant. We conclude that giving evaluators a marginally less selective garbling, and therefore (Lemma 1.14) any garbling that is less selective than S^G, cannot improve their payoffs. Likewise, for a marginally more selective garbling we have:

$$\lim_{\delta\to 0} \Pi\left(\sigma^G_{+\delta}; \mathcal{E}^G_{-\delta}\right) - \Pi\left(\sigma^G; \mathcal{E}^G\right) \propto -\lim_{\delta\to 0} \frac{\rho}{1-\rho} \times \frac{p_H(s^*)}{p_L(s^*)} \times \left(\frac{r^G_{H;+\delta}}{r^G_{L;+\delta}}\right)^{n-1} - \frac{c}{1-c} \geq 0$$

where the term on the RHS is now negative because the probability of trade *decreases* when strategies become more selective. By a reasoning similar to that behind Lemma 1.14, this

reveals that *no* garbling that is more selective can improve efficiency either.

This also proves that among those with optimal recommendations, the least selective garbling under which adverse selection is irrelevant cannot be improved with a more selective garbling and the most selective garbling under which adverse selection is not irrelevant cannot be improved with a less selective garbling.

□

References

Agarwal, Sumit et al. (July 2024). "Searching for Approval". In: *Econometrica* 92.4, pp. 1195–1231. DOI: 10.3982/ECTA18554.

Akerlof, George A. (1970). "The Market for "Lemons": Quality Uncertainty and the Market Mechanism". In: *The Quarterly Journal of Economics* 84.3, pp. 488–500. ISSN: 00335533, 15314650. DOI: 10.2307/1879431.

BCBS (Dec. 2017). *High-level summary of Basel III reforms*. Accessed: 2024-08-07. Bank for International Settlements. URL: https://www.bis.org/bcbs/publ/d424_hlsummary.pdf.

Bessembinder, Hendrik and William Maxwell (2008). "Markets: Transparency and the Corporate Bond Market". eng. In: *The Journal of Economic Perspectives* 22.2, pp. 217–234. ISSN: 0895-3309. DOI: 10.1257/jep.22.2.217.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades". In: *Journal of Political Economy* 100.5, pp. 992–1026. ISSN: 00223808, 1537534X. DOI: 10.1086/261849.

Blackwell, David (1953). "Equivalent Comparisons of Experiments". In: *The Annals of Mathematical Statistics* 24.2, pp. 265–272. ISSN: 00034851. DOI: 10.1214/aoms/1177729032.

Blackwell, David and Meyer A. Girshick (1954). *Theory of Games and Statistical Decisions*. eng. Wiley Publications in Statistics. New York ; London: John Wiley & Sons. ISBN: 0486638316.

Cavounidis, Costas (Sept. 2022). "Chunky Information and Adverse Selection". Working paper. URL: https://warwick.ac.uk/fac/soc/economics/staff/ccavounidis/chunky_information_and_adverse_selection.pdf (visited on 08/01/2025).

Chen, Kailin, Mehmet Ekmekci, and Stephan Lauermann (Feb. 2020). "Fishing for a Veto". Working paper. URL: https://drive.google.com/file/d/15fhv-YHDXkLdazeIdGwnxeNePxNKJsWc/view (visited on 08/01/2025).

Consumer Rights (2024). *How are credit checks conducted in the UK?* Accessed: 2024-10-26. URL: https://www.consumer-rights.org/finance/how-are-credit-checks-conducted-in-the-uk/.

Di Tillio, Alfredo, Marco Ottaviani, and Peter Norman Sørensen (2021). "Strategic Sample Selection". eng. In: *Econometrica* 89.2, pp. 911–953. ISSN: 0012-9682. DOI: 10.3982/ECTA17288.

Duffie, Darrell (2012). *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets*. Princeton University Press. ISBN: 9780691138961.

ECB (Feb. 2024). *ECB Guide to Internal Models*. European Central Bank. URL: https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.supervisory_guides202402_internalmodels.en.pdf (visited on 10/26/2024).

Financial Stability Board (Nov. 2017). *Artificial Intelligence and Machine Learning in Financial Services*. Financial Stability Board. URL: https://www.fsb.org/uploads/P011117.pdf (visited on 10/26/2024).

Gentzkow, Matthew and Emir Kamenica (2016). "A Rothschild-Stiglitz Approach to Bayesian Persuasion". In: *The American Economic Review* 106.5, pp. 597–601. ISSN: 00028282. DOI: 10.1257/aer.p20161049.

Glode, Vincent and Christian C Opp (June 2019). "Over-the-Counter versus Limit-Order Markets: The Role of Traders' Expertise". In: *The Review of Financial Studies* 33.2, pp. 866–915. ISSN: 0893-9454. DOI: 10.1093/rfs/hhz061.

Han, Lu and William C. Strange (2015). "The Microstructure of Housing Markets". eng. In: *Handbook of Regional & Urban Economics*. Vol. 5. Elsevier B.V, pp. 813–886. ISBN: 9780444595331. DOI: 10.1016/B978-0-444-59531-7.00013-2.

Herrera, Helios and Johannes Hörner (2013). "Biased social learning". In: *Games and Economic Behavior* 80.C, pp. 131–146. DOI: 10.1016/j.geb.2012.12.006.

Kay, John (2024). "HBOS Collapse Shows Danger of the Winner's Curse". In: *Financial Times*. URL: https://www.ft.com/content/f9254180-9200-11e5-bd82-c1fb87bef7af (visited on 10/26/2024).

Kreyszig, Erwin (1978). *Introductory Functional Analysis with Applications*. John Wiley & Sons.

Lauermann, Stephan and Asher Wolinsky (2016). "Search With Adverse Selection". In: *Econometrica* 84.1, pp. 243–315. ISSN: 00129682, 14680262. DOI: 10.3982/ECTA9969.

Lauermann, Stephan and Asher Wolinsky (2017). "Bidder Solicitation, Adverse Selection, and the Failure of Competition". In: *The American Economic Review* 107.6, pp. 1399–1429. ISSN: 00028282. DOI: 10.1257/aer.20131057.

Levin, Jonathan (2001). "Information and the Market for Lemons". In: *The RAND Journal of Economics* 32.4, pp. 657–666. ISSN: 07416261. URL: https://www.jstor.org/stable/2696386.

Milgrom, Paul R. (1979). "A Convergence Theorem for Competitive Bidding with Differential Information". In: *Econometrica* 47.3, pp. 679–688. ISSN: 00129682, 14680262. DOI: 10.2307/1910414.

Muirhead, R. F. (1900). "Inequalities relating to some Algebraic Means". In: *Proceedings of the Edinburgh Mathematical Society* 19, pp. 36–45. DOI: 10.1017/S0013091500032594.

Müller, Alfred and Marco Scarsini (2001). "Stochastic Comparison of Random Vectors with a Common Copula". eng. In: *Mathematics of operations research* 26.4, pp. 723–740. ISSN: 0364-765X.

Müller, Alfred and Dietrich Stoyan (n.d.). *Comparison Methods for Stochastic Models and Risks*. Wiley Series in Probability and Statistics. Wiley. ISBN: 9780471494461.

Munkres, J.R. (2000). *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated. ISBN: 9780131816299.

Rasmusen, Eric and Emmanuel Petrakis (1992). "Defining the Mean-Preserving Spread: 3-PT Versus 4-PT". In: *Decision Making Under Risk and Uncertainty: New Models and Empirical Findings*. Ed. by John Geweke. Dordrecht: Springer Netherlands, pp. 53–58. ISBN: 978-94-011-2838-4. DOI: 10.1007/978-94-011-2838-4_7.

Raymond, Lindsey (July 2024). "The Market Effects of Algorithms". Working paper. Accessed: 2024-10-26. URL: https://www.dropbox.com/scl/fi/22p85oogcf67mour5q8y2/LRaymond_JMP.pdf?rlkey=3v7nt884tx8y4rbxwgwvwhuu5&dl=0.

RE/MAX (2024). *Relisting at a Higher Price: Does it Work?* Accessed: 2024-10-26. URL: https://blog.remax.ca/relisting-at-a-higher-price-does-it-work/.

Riordan, Michael H. (1993). "Competition and bank performance: a theoretical perspective". In: *Capital Markets and Financial Intermediation*. Ed. by Colin Mayer and XavierEditors Vives. Cambridge University Press, pp. 328–343.

Rothschild, Michael and Joseph E Stiglitz (1970). "Increasing risk: I. A definition". In: *Journal of Economic Theory* 2.3, pp. 225–243. ISSN: 0022-0531. DOI: 10.1016/0022-0531(70)90038-4.

Shaked, M. and J.G. Shanthikumar (2007). *Stochastic Orders*. Springer Series in Statistics. Springer New York. ISBN: 9780387346755.

Stellantis Financial Services Italia S.p.A. (June 2024). *Auto ABS Italian Stella Loans S.r.l. Prospectus*. Auto ABS Italian Stella Loans S.r.l. URL: %7Bhttps://pcsmarket.org/wp-content/uploads/PROJECT-SIRIO-Prospectus-26.06-3.pdf%7D (visited on 10/30/2024).

Van Gestel, Tony and Bart Baesens (Oct. 2008). "Credit scoring". In: *Credit Risk Management: Basic Concepts: Financial Risk Components, Rating Analysis, Models, Economic and Regulatory Capital*. Oxford University Press. ISBN: 9780199545117. DOI: 10.1093/acprof:oso/9780199545117.003.0002.

Wilson, Robert (1977). "A Bidding Model of Perfect Competition". In: *The Review of Economic Studies* 44.3, pp. 511–518. ISSN: 00346527, 1467937X. DOI: 10.2307/2296904.

Wolinsky, Asher (1990). "Information Revelation in a Market with Pairwise Meetings". In: *Econometrica* 58.1, pp. 1–23. ISSN: 00129682, 14680262. DOI: 10.2307/2938332.

Zhu, Haoxiang (Apr. 2012). "Finding a Good Price in Opaque Over-the-Counter Markets". In: *The Review of Financial Studies* 25.4, pp. 1255–1285. ISSN: 0893-9454. DOI: 10.1093/rfs/hhr140.

Chapter 2

Optimally Dictatorial Committees

Abstract

I study the optimal voting mechanism for a committee that must decide whether to enact or block a policy of unknown benefit. Information can come both from committee members who can acquire it at cost, and a strategic lobbyist who wishes the policy to be enacted. I show that the *dictatorship of the most-demanding member* is a *dominant* voting mechanism: any other voting mechanism is (i) less likely to enact a good policy, (ii) more likely to enact a bad policy, and (iii) burdens each member with a greater cost of acquiring information.

*This chapter is version 2 of arXiv working paper 2507.21699 (https://arxiv.org/abs/2507.21699).

2.1 Introduction

[Democracy] is about making committees work.

Alan Bullock[1]

And a committee works only if it can harness information about the likely consequences of its decisions. This is a complex undertaking. The committee must first *obtain* the information it needs. To do so, it must fight on two fronts: at home, it must motivate its members to *acquire* information, even if this costs them time and effort. Outside, it must motivate external stakeholders to *provide* information, even if this may harm their interests. The committee must then effectively *aggregate* this information it has obtained: its decision must align with what this information warrants. How should such a committee be designed, then, to best achieve these—possibly competing—objectives?

I study this question in a canonical and parsimonious setting. There is a policy whose public benefit is unknown, and a committee wishes to decide whether to enact or block it. Committee members agree that a beneficial (*good*) policy should be enacted, and that a harmful (*bad*) one should be blocked. Nonetheless, each has different biases against the policy, so their preferred actions may differ while the policy's benefit remains uncertain. Outside the committee, there is a lobbyist. He prefers the policy to be enacted, whether it is good or bad.

To influence the committee, the lobbyist can provide information about the policy's benefit. He does so by designing a public experiment about it, at no cost. Subsequently, committee members can acquire additional information about the policy's benefit. Each member does so privately, and at personal cost. Finally, committee members vote. They cast their votes simultaneously, and a pre-determined *voting mechanism* prescribes (a) the set of votes each member may cast, and (b) the probability with which the policy will be enacted given the votes they cast.

I ask: is there a voting mechanism which secures

1. the lowest expected cost of information acquisition for each member,
2. lowest chance that a good policy is blocked,

[1] Frankland (2004).

3. and the lowest chance that a bad policy is enacted

among *all* possible voting mechanisms? Briefly put, is there a *dominant* voting mechanism?

In Theorem 2.1, I prove that a dominant voting mechanism does indeed exist by characterising one: *the dictatorship of the most-demanding member*. In it, only one member's vote influences the outcome. She is the *most-demanding member*: the one who would obtain the most certainty that the policy is good before enacting it, were she deciding without a committee or lobbyist to interfere.

There are a myriad of complex mechanisms. Why does this simple one dominate them all?

1) In a dictatorship, no member expects to acquire any information.

Clearly, the dictator's peers have no incentive to acquire any information—their vote never affects the outcome. The dictator, in contrast, would wish to acquire information unless she were sufficiently certain about the policy's unknown benefit. Crucially, the lobbyist always provides enough information to deliver that certainty. He sees no harm in doing so: if he did not, the dictator would simply acquire that information herself. He reaps ample benefit from this, too—I discuss this next. Thus, the dictator receives all her information from the lobbyist, and simply votes accordingly.

2) In a dictatorship, a good policy is never blocked.

A dictatorship incentivises the lobbyist to be the dictator's sole source of information—he provides enough information for the dictator to forgo acquiring any herself. The lobbyist leverages this position: he provides any additional information that raises the probability with which the policy is enacted. As such, he ensures that the dictator never blocks the policy unless she is certain that it is bad. Otherwise, he could provide her additional information that would either merely reinforce the dictator's belief that the policy should be blocked, or convince her that the policy is sufficiently likely to be good and turn her verdict around.

3) In the dictatorship of the most-demanding member, a bad policy is least likely to be enacted.

To enact the policy without acquiring further information, a dictator must be sufficiently certain that it is good. That threshold level of certainty is highest for the most-

demanding member: at her threshold, any dictator would enact the policy. I show that, in fact, *any* committee would enact the policy at this threshold—no committee is more demanding than the most-demanding member. Intuitively, a committee member values information the most when she is a dictator; other voting mechanisms might anyway fail to incorporate her information into the final outcome. So, at a belief threshold where no dictator would acquire information before enacting the policy, no committee member would do so either before voting to enact it.

As no committee is more demanding than the most-demanding member, the lobbyist can ensure that no committee is less likely to enact the policy than the most-demanding member's dictatorship. But the most-demanding member's dictatorship always enacts good policies; so, I conclude that no committee is less likely to enact a bad policy than the most-demanding member's dictatorship.

Contribution and Literature

I study how a committee must be designed in the face of its fourfold informational problem:

1. it must incentivise external stakeholders to provide information relevant to its decision,
2. it must incentivise its members to acquire such information, even if they find this costly,
3. it must minimise the cost its members bear as they acquire such information,
4. its final decision must efficiently aggregate the information it has accumulated.

There is an extensive and multidisciplinary body of literature around this fourfold informational problem. However, existing literature addresses only a subset of these four problems at a time, ignoring the rest. An earlier strand studies environments without external stakeholders to interfere with the committee's decision. This strand asks how a committee may aggregate its members' private information—both when information is held exogenously[2] and when it is acquired at personal cost.[3] A more recent strand, in turn,

[2]See Condorcet (1785)'s celebrated Jury Theorem, and the seminal works of Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996), (1997, and 1998.

[3]See Gersbach (1995), Li (2001), Persico (2004), Smorodinsky and Tennenholtz (2004) and (2006), Gerardi and Yariv (2008), Gershkov and Szentes (2009), and Chan, Lizzeri, et al. (2018). Stephenson (2011) offers a legal-theoretic discussion and treatment.

focuses on external stakeholders who may try to influence the committee through providing information, but mutes committee members' ability to acquire additional information.[4]

However, the objectives these four problems impose on the committee are often in conflict. Consequently, so is the guidance offered by these two strands of the literature on optimal committee design.[5] This is troubling for policymakers who may want to apply these lessons, since policymaking committees often face all four of these problems simultaneously.[6]

This paper addresses this gap by studying the optimal voting mechanism for a committee facing all four of these problems simultaneously. In Theorem 2.1, I find that a dictatorship dominates any other voting mechanism. This is robust to committee members' information acquisition costs and biases towards the policy; these characteristics only determine *who* the dictator should be. This robustness contrasts with the earlier literature on information acquisition in committees, which generally finds that the optimal mechanism may depend greatly on committee members' characteristics.[7]

Among papers in the aforementioned literature, Theorem 2.1 aligns most closely with Alonso and Câmara (2016) and Bardhi and Guo (2018). In both, a politician, or Sender (replacing *the lobbyist* here), provides information to persuade a committee whose members have no access to private information. Both find that, when members' preferences align to the extent they do in this paper,[8] each member prefers the unanimity rule to any other k-

[4]See Breyer (1982) and Coglianese, Zeckhauser, and Parson (2004) for legal-theoretic discussions. See Alonso and Câmara (2016), Bardhi and Guo (2018), Chan, Gupta, et al. (2019), and Heese and Lauermann (forthcoming) for contributions to Economics.

[5]For instance, both Persico (2004) and Alonso and Câmara (2016) study the optimal k-majority rule in the (suitably adapted) canonical voting setting this paper examines, preceded by Austen-Smith and Banks (1996). In Persico (2004), there are no external stakeholders to inform the committee, unlike Alonso and Câmara (2016). In contrast, Alonso and Câmara (2016)'s voters (replacing the *committee members* here) have no access to private information, unlike Persico (2004). Consequently, Persico (2004) finds that large majority rules are often suboptimal, as they disincentivise committee members from acquiring private information. In contrast, Alonso and Câmara (2016) find that a unanimity rule is optimal when voters' preferences are sufficiently aligned, as this rule maximises the politician's (replacing *the lobbyist* here) incentives to provide information to the voters.

[6]Breyer (1982) highlights this for agencies regulating public utilities: "[An agency] will obtain the information, in part, through research by agency staff, as they consult research literature and talk to employees of other agencies. ... Once the Notice of Proposed Rulemaking is promulgated, however, ... information is then more likely to flow in the form of comments submitted to the agency by lawyers for parties.

[7]For instance, Persico (2004) finds that the optimal majority rule depends on the precision of the information available to committee members. Gerardi and Yariv (2008) find that the optimal size of a committee depends on its members' cost of acquiring information, and how precise that information is. Gershkov and Szentes (2009) find that the whether the optimal sequential voting rule is ex-post efficient similarly depends on agents' (replacing the *committee members* here) cost of acquiring information.

[8]Both consider a more general class of preferences for members than this paper. Alonso and Câmara

majority rule.[9] This is because under unanimity, the *most-demanding member*—there, the one with the strongest bias against the policy—becomes the pivotal voter. The politician, or Sender, must then persuade her to enact the policy. This compels him to provide more information to the committee than he would under any other k-majority rule.

I also find that under the dictatorship of the most-demanding member, the lobbyist is incentivised to provide more information to the committee than under any other voting mechanism (Proposition 2.2). However, this only provides a partial intuition for Theorem 2.1: even if such a voting mechanism extracts less information from the lobbyist, it may still provide strong information acquisition incentives to committee members. As such, such a voting mechanism may still lead to better outcomes than the dictatorship of the most-demanding member—in principle.

What drives Theorem 2.1 is that the lobbyist *always* prevents a voting mechanism from outperforming the dictatorship of the most-demanding member as such. When the most-demanding member is the dictator, the information she receives from the lobbyist dissuades her from acquiring more herself. In proving Theorem 2.1, I establish that the same amount of information dissuades *any* committee member from acquiring information, regardless of the voting mechanism. I then show that the lobbyist provides less than this amount of information only when this persuades the committee to enact bad policies more often—potentially, even at the cost of enacting good policies less often.

Matysková and Montes (2023) study a model akin to mine, but where the Sender (here, the *lobbyist*) attempts to persuade a single decision-maker (here, a *committee member*) rather than a committee. Their main result finds that it is without loss to restrict to equilibria where the Sender gives the Receiver enough information to dissuade her from acquiring any herself. My Lemma 2.1, which finds that the lobbyist always provides enough information to disincentivise the dictator from acquiring any herself, can be viewed as a manifestation of their main result.

Lastly, it is worth remarking that this paper only studies *simultaneous* voting mecha-

(2016) find this result when members agree on which states make enacting the policy more desirable, and in which states it should be enacted. Bardhi and Guo (2018) find it when members' preference states are perfectly correlated.

[9]Bardhi and Guo (2018) study both the case where information provision is *general*—the Sender publicly communicates all information—and the case where it is *private*—the Sender can provide different information to different members. They find this for the former, but not necessarily the latter case.

nisms. Some authors[10] have also studied what the optimal *sequential* voting mechanism is for a committee. Investigating this in settings where the committee must incentivise both information acquisition by its members and information provision by a strategic lobbyist remains a promising area of research.

Outline

The remainder of the paper is organised as follows. Section 2.2 introduces the model. Section 2.3 introduces dictatorships and presents Lemma 2.1, which discusses the equilibria dictatorships induce. Section 2.4 presents and proves Theorem 2.1, the main result. It also presents Proposition 2.2, which proves that no voting mechanism incentivises the lobbyist to provide more information than he does under the dictatorship of the most-demanding member. Section 2.5 concludes by discussing how sensitive the results of the paper are to certain modelling assumptions.

2.2 The Model

A committee of N *members*, $m_1, ..., m_N$, wish to decide whether to *enact* or *block* a policy. The benefit member m_i obtains from enacting the policy depends on an unknown state of the world, $\omega \in \{\omega_0, \omega_1\}$. When $\omega = \omega_1$—the policy is *good*—she obtains a payoff of $u_i \geq 0$ from enacting it. When $\omega = \omega_0$—the policy is *bad*—she incurs a payoff loss of 1 instead. Blocking the policy, in contrast, yields her a payoff of 0 in either state. There is also a *lobbyist* outside the committee. He prefers the policy to be enacted regardless of the state.

At the outset of the game, the lobbyist and committee members share the prior belief that $\omega = \omega_1$ with probability $p \in (0, 1)$. The lobbyist can then provide additional information about the state. He does so by committing to perform and publicly disclose the outcome of an (Blackwell) experiment of his choosing. Formally, an experiment is comprised of a measurable space $(X, \mathcal{X})$ and, for each possible state ω_k, a probability measure P_k over this space. Upon observing the outcome of the lobbyist's experiment, players use Bayes Rule to update their prior belief p to an *interim belief* $r \in [0, 1]$.

I follow a *belief-based approach*[11] and model the lobbyist as directly choosing the prob-

[10] For instance, Gershkov and Szentes (2009), and Chan, Lizzeri, et al. (2018).
[11] See, for instance, Kamenica and Gentzkow (2011), and Caplin and Dean (2013).

ability measure μ_l his experiment will induce over the space of interim beliefs, $[0,1]$. He can choose any such measure, provided it is *Bayes plausible*; i.e. $\int_{[0,1]} r\mu_l(dr) = p$. I denote the set of such measures as $\Delta_p([0,1])$.

After the lobbyist reveals the outcome of his experiment, each committee member may acquire additional information about the state. She may do so by choosing an additional experiment whose outcome she observes. Information acquisition by committee members is *private*; neither m_i's chosen experiment nor its outcome are disclosed to anyone.

Following the outcome of her experiment, m_i updates her interim belief r to a *posterior belief* $q_i \in [0,1]$ that $\omega = \omega_1$. Again, I model each member m_i as choosing the probability measure $\mu_i(. \mid r) \in \Delta_r([0,1])$ her experiment will induce over her posterior beliefs, given her interim belief r. To simplify exposition, I assume that the outcomes of members' chosen experiments are independent conditional on the true state ω.

Producing and providing public information is costless for the lobbyist. In contrast, acquiring private information is costly for committee members: m_i must pay a cost $C_i(\mu)$ to observe the outcome of an experiment which induces the distribution μ over her possible posterior beliefs. The cost function C_i maps $\Delta([0,1])$, the set of all probability measures over $[0,1]$, to $\mathbb{R}^+$. It is *uniformly posterior separable*;[12] i.e., there is a convex function $c_i : [0,1] \to \mathbb{R}^+$ such that:

$$C_i(\mu) = \int_{[0,1]} c_i(q)d\mu(q) - c_i\left(\mathbb{E}_\mu[q]\right)$$

Finally, each committee member casts a private vote into a *voting mechanism*. A voting mechanism is a tuple $\left(D, \{V_i\}_{i=1}^N, \{\mathcal{V}_i\}_{i=1}^N\right)$ where $V_i \ni v_i$ is a set of possible votes each member m_i may cast, $\mathcal{V}_i$ is a σ-algebra over V_i, and the *decision rule* $D : \bigtimes_{i=1}^N V_i \to [0,1]$ is a measurable mapping from the space of possible vote profiles to the probability that the policy is enacted. I require that:

1. given others' votes, member m_i have a *lowest* and *highest* vote to submit:[13]

[12] See, for instance, Caplin, Dean, and Leahy (2017), and Denti (2022).

[13] This requirement is automatically satisfied when the set of possible vote profiles $\bigtimes_{i=1}^N V_i$ is finite, or when $\bigtimes_{i=1}^N V_i$ is compact and the decision rule $D(\cdot)$ is continuous.

$$\arg\min_{v_i \in V_i} D(v_i, \mathbf{v_{-i}}) \text{ and } \arg\max_{v_i \in V_i} D(v_i, \mathbf{v_{-i}}) \text{ exist for each } \mathbf{v_{-i}} \in \mathbf{V_{-i}} := \bigtimes_{j \neq i} V_j$$

2. there be a profile of votes upon which the policy is almost surely enacted (blocked):

$$\exists\ \underline{\mathbf{v}}, \bar{\mathbf{v}} \in \bigtimes_{i=1}^{N} V_i \text{ such that } D(\underline{\mathbf{v}}) = 0 \text{ and } D(\bar{\mathbf{v}}) = 1$$

For any given voting mechanism, I focus on the lobbyist-preferred perfect Bayesian equilibrium of this game. A perfect Bayesian equilibrium consists of:

1. a voting strategy $\{\tau_i^*(. \mid q, r)\}_{(q,r)\in[0,1]^2}$ for each member m_i, where $\tau_i^*(. \mid q, r) \in \Delta(V_i)$,

2. an information strategy $\{\mu_i^*(. \mid r)\}_{r\in[0,1]}$ for each member m_i, where $\mu_i^*(. \mid r) \in \Delta_r([0,1])$,

3. a strategy $\mu_l^* \in \Delta_p([0,1])$ for the lobbyist,

such that:

1. Given other members' strategies, member m_i's voting strategy maximises her expected payoff at any pair of interim and posterior beliefs $(r, q) \in [0,1]^2$ she may hold:

$$\begin{aligned} & v_i \in \operatorname{supp} \tau_i^*(. \mid q, r) \\ \implies & v_i \in \operatorname*{argmax}_{v \in V_i} \int_{[0,1]^{N-1}} U_i(r, q_i, \mathbf{q_{-i}}) \times \left\{ \int_{V_{-i}} D(v, \mathbf{v_{-i}}) \tau_{-i}^*(d\mathbf{v_{-i}} \mid r, \mathbf{q_{-i}}) \right\} \mu_{\mathbf{-i}}^*(d\mathbf{q_{-i}} \mid r, q_i) \end{aligned}$$

where:

i. $\mu_{\mathbf{-i}}^*(\cdot \mid r, q_i)$ is the probability measure m_i assigns over the posterior beliefs of all members excluding herself, conditional on her interim and posterior beliefs r and q_i, and those members' information strategies.[14]

ii. $\tau_{-i}^*(\cdot \mid r, \mathbf{q}_{-i}) := \bigtimes_{j \neq i} \tau_j^*(. \mid r, q_j)$ is the probability measure over the votes of all members excluding m_i, conditional on their interim and posterior beliefs r and $\mathbf{q}_{-i}$, and their voting strategies.

[14] This measure can easily be derived using Bayes' Rule.

iii. $U_i(r, q_i, \mathbf{q_{-i}})$ denotes the payoff m_i expects from enacting the policy given all players' information about the state:

$$U_i(r, q_i, \mathbf{q_{-i}}) := \mathbb{P}_{r,q_1,\ldots q_N}(\omega_1) \times (u_i + 1) - 1$$

Here, $\mathbb{P}_{r,q_1,\ldots q_N}(\omega_1)$ denotes the probability that $\omega = \omega_1$ given the same information:

$$\mathbb{P}_{r,q_1,\ldots,q_N}(\omega_1) := \left(\frac{r}{1-r} \times \prod_{j=1}^{N} \frac{q_j/(1-q_j)}{r/(1-r)} \right) \times \left(1 + \frac{r}{1-r} \times \prod_{j=1}^{N} \frac{q_j/(1-q_j)}{r/(1-r)} \right)^{-1}$$

2. Given others' strategies, member m_i's information strategy maximises her expected payoff at any interim belief r:

$$\mu_i^*(\cdot \mid r) \in \underset{\mu \in \Delta_r([0,1])}{\arg\max} \int \pi_i^*(r, q)\mu(dq) - C_i(\mu)$$

where $\pi_i^*(r, q)$ is the payoff m_i expects in the ensuing subgame, conditional on her interim belief r, posterior belief q, and ensuing voting strategy,

3. The lobbyist's information strategy μ_l^* maximises the probability that the policy is enacted given committee members' information and voting strategies.

The lobbyist-preferred perfect Bayesian equilibrium (hereon, simply *equilibrium*) is the perfect Bayesian equilibrium where the probability that the committee enacts the policy is highest.

Dominant Voting Mechanisms

A voting mechanism is *dominant* if it gives rise to an equilibrium that yields:

i. a (weakly) higher probability that the policy is blocked when $\omega = \omega_0$,

ii. a (weakly) higher probability that the policy is enacted when $\omega = \omega_1$, and

iii. a (weakly) lower expected cost of information acquisition for each committee member

compared to any equilibrium that any other voting mechanism yields. Every committee member (weakly) prefers a dominant voting mechanism over any other voting mechanism.

2.3 Dictatorships

One of the simplest voting mechanisms the committee may adopt is a *dictatorship*. In a dictatorship, only one member's vote influences the outcome.

Definition 2.1. A voting mechanism $(\overset{\bigstar}{D}_i, \{V_j\}_{j=1}^N, \{\mathcal{V}_j\}_{j=1}^N)$ is an m_i-*dictatorship* if:

$$\overset{\bigstar}{D}_i(v_i, \mathbf{v}_{-i}) = \overset{\bigstar}{D}_i(v_i, \mathbf{v}'_{-i}) \quad \text{for any } v_i \in V_i \text{ and all } \mathbf{v}_{-i}, \mathbf{v}'_{-i} \in \mathbf{V_{-i}}$$

Lemma 2.1 characterises the equilibria dictatorships induce.

Lemma 2.1. In the equilibrium of an m_i-dictatorship:

1. m_i's voting strategy satisfies:

$$\mathbb{E}_{\tau_i^*(.|r,q)} \overset{\bigstar}{D}_i = \begin{cases} 0 & \text{for } q < (1+u_i)^{-1} \\ 1 & \text{otherwise} \end{cases}$$

   for any interim belief $r \in [0,1]$.

2. m_i's information acquisition strategy $\{\mu_i^*(. \mid r)\}_{r\in[0,1]}$ is the unique one such that for every $r \in [0,1]$, $\mu_i^*(\cdot \mid r) \in \Delta_r([0,1])$, and:

$$\text{supp } \mu_i^*(\cdot \mid r) = \begin{cases} \left\{\underline{q}_i^*, \bar{q}_i{}^*\right\} & r \in \left(\underline{q}_i^*, \bar{q}_i{}^*\right) \\ \{r\} & \text{otherwise} \end{cases}$$

   where $\underline{q}_i^*$ and $\bar{q}_i{}^*$ are two threshold beliefs such that $\underline{q}_i^* \leq (1+u_i)^{-1} \leq \bar{q}_i{}^*$.

3. The lobbyist's information strategy μ_l^* is the unique probability measure satisfying:

$$\text{supp } \mu_l^* = \begin{cases} \{0, \bar{q}_i{}^*\} & p < \bar{q}_i{}^* \\ \{p\} & \text{otherwise} \end{cases}$$

4. Members except m_i never incur a positive information acquisition cost; i.e., $C_j\left(\mu_j^*(. \mid r)\right) = 0$ for every $r \in [0,1]$.

In an m_i-dictatorship, m_i has exclusive control over the decision. Thus, she votes to

ensure that the decision maximises her expected payoff: the policy is enacted whenever she expects a weakly positive payoff from doing so, and it is blocked otherwise.

m_i's information strategy is fully characterised by her two *thresholds of persuasion*, $\underline{q}_i^*$ and $\bar{q}_i^*$. m_i acquires additional information about the state if and only if her interim belief r falls between these thresholds. When she does, the information she acquires moves her posterior belief precisely to one of these two thresholds. Importantly, m_i's thresholds of persuasion always bracket her *threshold of indifference* $(1+u_i)^{-1}$; i.e., the belief which leaves her indifferent between enacting and blocking the policy. This reflects that m_i never acquires a costly experiment whose outcome always leads to the same vote, and as such, has no value for her decision. In contrast, m_i acquires no further information when her interim belief lies beyond these thresholds, $r \notin (\underline{q}_i^*, \bar{q}_i{}^*)$. Instead, she simply takes whatever decision her interim belief favours. Her peers, having no influence over the outcome, never acquire costly information about the state.

In turn, the lobbyist wishes to maximise the probability that the committee enacts the policy. In an m_i-dictatorship, this amounts to maximising the probability that m_i's posterior belief weakly exceeds her threshold of indifference, $(1+u_i)^{-1}$, so that she votes for enactment. Owing to m_i's simple information strategy, the lobbyist can restrict himself to interim beliefs beyond m_i's thresholds of persuasion, $r \notin \left(\underline{q}_i^*, \bar{q}_i^*\right)$. Interim beliefs between these thresholds merely prompt m_i to distribute her posterior beliefs back to her thresholds of persuasion. In contrast, those outside that interval prompt her to vote without acquiring further information.[15]

This reduces the lobbyist's problem to a simple Bayesian Persuasion exercise a là Kamenica and Gentzkow (2011). When m_i's prior belief already exceeds her upper threshold of persuasion $\bar{q}_i^*$, the lobbyist provides no further information. Then, m_i simply enacts the policy without acquiring any additional information. Otherwise, the lobbyist's experiment either induces the interim belief 0, or $\bar{q}_i^*$. The latter precisely equals m_i's upper threshold of persuasion: m_i is persuaded to enact the policy, but with no more evidence than necessary. The former, on the other hand, conclusively reveals that the policy is bad, persuading

[15] Matysková and Montes (2023) establish a more general version of this argument. A Sender (replacing *the lobbyist* here) who faces a single Receiver (replacing *the dictator* here) with a uniformly posterior separable information acquisition cost suffers no loss from restricting himself to interim beliefs where the Receiver acquires no further information.

m_i to block the policy. Only a bad policy can induce this interim belief, ensuring that m_i never blocks a good policy.

Proof, Lemma 2.1:

1. Since only m_i"s vote affects the outcome, we can define the measurable function $\overset{\blacktriangledown}{f}_i : V_i \to [0,1]$ to be $\overset{\blacktriangledown}{f}_i(v_i) := \overset{\blacktriangledown}{D}_i(v_i, \mathbf{v_{-i}})$ for an arbitrary $\mathbf{v_{-i}} \in V_{-i}$. Then, $v_i \in \text{supp } \tau_i^*(\cdot \mid r, q)$ only if:

$$v_i \in \arg\max_{v \in V_i} \int_{[0,1]^{N-1}} U_i(r, q_i, \mathbf{q_{-i}}) \times \left\{ \int_{V_{-i}} \overset{\blacktriangledown}{f}_i(v_i) \tau_{-i}^*(d\mathbf{v_{-i}} \mid r, \mathbf{q_{-i}}) \right\} \mu_{\mathbf{-i}}^* (d\mathbf{q_{-i}} \mid r, q_i)$$

$$\implies v_i \in \arg\max_{v \in V_i} \overset{\blacktriangledown}{f}_i(v_i) \times \int_{[0,1]^{N-1}} U_i(r, q_i, \mathbf{q_{-i}}) \mu_{\mathbf{-i}}^* (d\mathbf{q_{-i}} \mid r, q_i)$$

$$\implies v_i \in \arg\max_{v \in V_i} \overset{\blacktriangledown}{f}_i(v_i) \times [q_i \times (1 + u_i) - 1]$$

where the last implication follows from the martingale property of beliefs; i.e., $\mathbb{E}_{\mathbf{q_{-i}}}[\mathbb{P}_{r,q_1,\ldots,q_N} \mid q_i] = q_i$.

Thus, $\{\tau_i^*(\cdot \mid r, q)\}_{(r,q) \in [0,1]^2}$ is an equilibrium strategy only if:

$$\mathbb{E}_{\tau_i^*(\cdot|r,q)} \overset{\blacktriangledown}{f}_i(.) = \begin{cases} 0 & q < (1+u_i)^{-1} \\ 1 & q > (1+u_i)^{-1} \end{cases}$$

When $q = (1+u_i)^{-1}$, m_i is indifferent between enacting and blocking the policy, hence any vote. He must then submit the lobbyist-preferred vote; i.e., $\mathbb{E}_{\tau_i^*(\cdot|r,(1+u_i)^{-1})} \overset{\blacktriangledown}{f}_i(.) = 1$.

2. m_i only needs to decide whether to vote to maximise or minimise the enactment probability of the policy. So, she can restrict to information strategies where $|\text{supp } \mu_i^*(\cdot \mid r)| \leq 2$, without loss of generality. Furthermore, m_i can always restrict herself to either acquiring no information, or acquiring information that has decision-making value; i.e., $(1+u_i)^{-1} \in \text{conv}\,(\text{supp } \mu_i^*(\cdot \mid r))$ whenever $|\text{supp } \mu_i^*(\cdot \mid r)| = 2$. Lastly, note that Theorem 1 in Caplin, Dean, and Leahy (2022) implies that m_i's strategy must be *Locally Posterior Invariant* when C_i is uniformly posterior separa-

ble; i.e., for any $\tilde{r} \in \text{conv}\,(\text{supp}\ \mu_i^*(\cdot \mid r))$, $\mu_i^*(\cdot \mid \tilde{r})$ is the unique distribution such that $\text{supp}\ \mu_i^*(\cdot \mid \tilde{r}) = \text{supp}\ \mu_i^*(\cdot \mid r)$.

Letting $\left\{\underline{q}_i^*, \bar{q}_i^*\right\} := \text{supp}\ \mu_i^*(\cdot \mid (1+u_i)^{-1})$ and $\underline{q}_i^* \leq \bar{q}_i^*$, we thus conclude:

(a) for all $r \in \left[\underline{q}_i^*, \bar{q}_i^*\right]$, $\mu_i^*(\cdot \mid r)$ is the unique distribution such that $\text{supp}\mu_i^*(\cdot \mid r) = \left\{\underline{q}_i^*, \bar{q}_i^*\right\}$.

(b) for all $r \notin \left[\underline{q}_i^*, \bar{q}_i^*\right]$, $\mu_{i;r}^* = \delta_r$.

3. Given members' voting and information strategies, let $\zeta_i^*(r)$ denote the probability that the policy is enacted in an m_i-dictatorship, following an interim belief of r. Denote the concave envelope of this function as $\hat{\zeta_i}^*(r)$. These two functions are given by:

$$\zeta_i^*(r) := \begin{cases} 0 & r \leq \underline{q}_i^* \\ \frac{r-\underline{q}_i^*}{\bar{q}_i^*-\underline{q}_i^*} & r \in \left[\underline{q}_i^*, \bar{q}_i^*\right] \\ 1 & r \geq \bar{q}_i^* \end{cases} \qquad \hat{\zeta_i}^*(r) := \begin{cases} \frac{r}{\bar{q}_i^*} & r \leq \bar{q}_i^* \\ 1 & r \geq \bar{q}_i^* \end{cases}$$

That $\hat{\zeta_i}^*(r)$ is indeed the concave envelope of $\zeta_i^*(r)$ can be easily verified: any concave function that coincides with ζ_i^* at $r = 0$ and $r = \bar{q}_i^*$ must weakly lie above $\hat{\zeta_i}^*(r)$ on the interval $[0, \bar{q}_i^*]$.

The condition $\mathbb{E}_\mu\, \zeta_i^*(.) = \hat{\zeta_i}^*(p)$ is then sufficient for the optimality of μ for the lobbyist (Kamenica and Gentzkow (2011)). It is easily seen that the distribution μ_l^* defined in Lemma 2.1 satisfies this condition.

□

2.4 A Dominant Dictatorship

The equilibrium of an m_i-dictatorship is shaped by the dictator's upper threshold of persuasion, $\bar{q_i}^*$. In this equilibrium, no committee member acquires any information; the committee relies entirely on the information provided by the lobbyist. In turn, the lobbyist ensures that the committee blocks the policy only when it is conclusively revealed to be bad. In contrast, he never *over-persuades* m_i, the dictator, to enact the policy; m_i's

interim belief never *strictly* exceeds her upper threshold of persuasion, $\bar{q}_i{}^*$, unless her prior belief already does.

So, the higher a dictator's upper threshold of persuasion is, the more information she extracts from the lobbyist. I call the member for whom this threshold is highest the *most-demanding member* of the committee.

Definition 2.2. Member m_i is the *most-demanding member* of the committee, denoted $\overset{\bigstar}{m}$, if $\bar{q}_i^* = \max\{\bar{q}_1^*, \bar{q}_2^*, ..., \bar{q}_N^*\} =: \overset{\bigstar}{q}$.

The dictatorship of the most-demanding member clearly dominates the dictatorship of any other member—no dictatorship results in costly information acquisition or a good policy getting blocked, but the $\overset{\bigstar}{m}$-dictatorship minimises the probability of a bad policy getting enacted. Theorem 2.1 shows that, in fact, the $\overset{\bigstar}{m}$-dictatorship dominates *every* voting mechanism.

Theorem 2.1. The dictatorship of the most-demanding member is a dominant voting mechanism.

Theorem 2.1 rests on two fundamental insights:

1. The $\overset{\bigstar}{m}$-dictatorship dominates any voting mechanism that is at least as likely to enact the policy as itself.

2. Any voting mechanism is at least as likely to enact the policy as the $\overset{\bigstar}{m}$-dictatorship.

The first is easy to see. The $\overset{\bigstar}{m}$-dictatorship clearly minimises committee members' information costs: they incur none. While doing so, it also ensures that a good policy is almost surely enacted. So, a voting mechanism which has a greater probability of enacting the policy than the $\overset{\bigstar}{m}$-dictatorship must enact bad policies more often. Such a mechanism is clearly dominated by the $\overset{\bigstar}{m}$-dictatorship.

To see the second, suppose the lobbyist ignores the committee's intricate voting mechanism and conducts the experiment tailored for the $\overset{\bigstar}{m}$-dictatorship. I show that, with this experiment, the committee enacts the policy no less often than the $\overset{\bigstar}{m}$-dictatorship would. Indeed, the $\overset{\bigstar}{m}$-dictatorship enacts the policy if and only if this experiment generates the interim belief $\overset{\bigstar}{q}$. I show that this interim belief persuades committee members to vote for

enactment *regardless* of the voting mechanism they use. This is because $\overset{\bigstar}{q}$ (weakly) exceeds any member's upper threshold of persuasion; if a member were deciding alone, she would enact the policy without acquiring further information. Whenever she believes her peers have no private information relevant to the decision anyway, she wishes to vote the same way.

Proof, Theorem 2.1: Fix $\left(\{V_j\}_{j=1}^N, \{\mathcal{V}_j\}_{j=1}^N, D\right)$ to be an arbitrary voting mechanism, and let $\overset{\bigstar}{D}$ denote the voting rule for the $\overset{\bigstar}{m}$-dictatorship.

I will use Claim 2.1 to establish Theorem 2.1:

Claim 2.1. Any voting mechanism induces an equilibrium where, in the subgames with interim beliefs $r \geq \overset{\bigstar}{q}$, committee members:

(a) acquire no information: their information strategies $\left\{\mu_j^*(\cdot \mid r)\right\}_{r\in[0,1]}$ are such that for all $r \geq \overset{\bigstar}{q}$, $\mu_j^*(\cdot \mid r) = \delta_r$.

(b) vote to enact the policy: their voting strategies $\left\{\tau_j^*(\cdot \mid r, q)\right\}_{(r,q)\in[0,1]^2}$ are such that for all $r \geq \overset{\bigstar}{q}$, $\tau_j^*(\cdot \mid r, r) = \delta_{v_j^*}$ where $D(v_1^*, ..., v_N^*) = 1$.

Proof, Claim 2.1: Clearly, no strategy profile can leave the lobbyist strictly better off in the relevant subgames. So, it suffices to show that the asserted strategies are a mutual best response for all committee members.

Let us first establish part (b). Suppose every member except m_i employs the asserted strategies. Define $\pi_i(r, q)$ to be the payoff m_i expects under this voting mechanism, given her interim belief r, posterior belief q, others' strategies, and her voting strategy. When $r \geq \overset{\bigstar}{q}$, other committee members acquire no information; so, their votes are not informative about the state. As such, for $r \geq \overset{\bigstar}{q}$:

$$\pi_i(r, q) = [q \times (1 + u_i) - 1] \times \int_{V_i} D(v, \mathbf{v}_{-\mathbf{i}}^*)\tau_i(dv)$$

given m_i's chosen voting strategy $\{\tau_i(\cdot \mid r, q)\}_{(r,q)\in[0,1]^2}$ and her peers' equilibrium votes, $\mathbf{v}_{-\mathbf{i}}^*$. When q_i is weakly above (strictly below) $(1 + u_i)^{-1}$, $\pi_i(r, q)$ is strictly increasing (decreasing) in the enactment probability. So, m_i must vote to maximise (minimise) this

probability. Since $\overset{\blacktriangledown}{q} \geq \bar{q}_i^* \geq (1+u_i)^{-1}$, this implies that m_i cannot profitably deviate from the asserted voting strategies. This proves part (b) of the claim.

Now, let us prove part (a). The argument above established that, in equilibrium:

$$\pi_i(r,q) = \begin{cases} q \times (1+u_i) - 1 & q \geq (1+u_i)^{-1} \\ \min\limits_{v_i \in V_i} D(v_i, \mathbf{v}_{-\mathbf{i}}^*) \times [q \times (1+u_i) - 1] & q < (1+u_i)^{-1} \end{cases}$$

for any $r \geq \overset{\blacktriangledown}{q}$. To prove that acquiring no information is a best response for m_i, we must show:

$$\int_{[0,1]} [\pi_i(q;r) - \pi_i(r;r)]\, d\mu(q) - C_i(\mu) \leq 0 \qquad \text{for all } r \geq \overset{\blacktriangledown}{q} \text{ and } \mu \in \Delta_r([0,1]) \tag{2.4.1}$$

To this end, define $\overset{\blacktriangledown}{\pi}_i(q)$ to be the payoff m_i expects from her preferred action at a posterior belief q; $\overset{\blacktriangledown}{\pi}_i(q) := \max\{0, q \times (1+u_i) - 1\}$. Since $\bar{q}_i^* \geq \overset{\blacktriangledown}{q}$, m_i acquires no information for interim beliefs $r \geq \overset{\blacktriangledown}{q}$ under her dictatorship. This implies:

$$\int_{[0,1]} \left[\overset{\blacktriangledown}{\pi}_i(q) - \overset{\blacktriangledown}{\pi}_i(r)\right] d\mu(q) - C_i(\mu) \leq 0 \qquad \text{for all } r \geq \overset{\blacktriangledown}{q} \text{ and } \mu \in \Delta_r([0,1]) \tag{2.4.2}$$

Now let $\tilde{\pi}_i(q,r) := [\pi_i(q,r) - \pi_i(r,r)] - [\overset{\blacktriangledown}{\pi}_i(q) - \overset{\blacktriangledown}{\pi}_i(r)]$. 2.4.1 follows from 2.4.2 once we prove:

$$\int_{[0,1]} \tilde{\pi}_i(q,r) d\mu(q) \leq 0 \qquad \text{for all } r \geq \overset{\blacktriangledown}{q} \text{ and } \mu \in \Delta_r([0,1]) \tag{2.4.3}$$

To see 2.4.3, first consider the behaviour of $\tilde{\pi}_i(q,r)$ as a function of q, for some fixed r. This function is piecewise linear, has a kink at $q = (1+u_i)^{-1}$, increases before this kink, and decreases after it. Therefore, for any fixed r, $\tilde{\pi}_i(q,r)$ is a concave function of q. Moreover, $\tilde{\pi}(r,r) = 0$, trivially. By Jensen's Inequality, we then obtain:

$$\int_{[0,1]} \tilde{\pi}_i(q;r) d\mu(q) \leq \tilde{\pi}_i(r;r) = 0 \qquad \text{for all } \mu \in \Delta_r([0,1])$$

□

In the equilibrium the $\overset{\bigstar}{m}$-dictatorship induces, no committee member expects to incur information acquisition costs. Moreover, a good policy is almost surely enacted. Therefore, to prove Theorem 2.1, it suffices to show that any other mechanism has a higher equilibrium probability of enacting a bad policy.

In the equilibrium of the $\overset{\bigstar}{m}$-dictatorship, the policy is enacted if and only if the lobbyist's experiment induces the interim belief $\overset{\bigstar}{q}$. Claim 2.1 establishes that this interim belief leads to enactment under any voting mechanism. So, by performing the same experiment, the lobbyist can secure a weakly greater probability of enactment under any such mechanism. This implies that any such mechanism is likelier to enact the policy than the $\overset{\bigstar}{m}$-dictatorship, in equilibrium. Since the $\overset{\bigstar}{m}$-dictatorship never blocks a bad policy, any such mechanism must be likelier to enact a bad policy, too. □

Theorem 2.1 aligns with the findings of Alonso and Câmara (2016) and Bardhi and Guo (2018). These papers also feature a lobbyist[16] who wishes to persuade the committee to enact a policy by designing an experiment about its unknown benefit. However, their committee members have no access to private information. They find that, when committee members' preferences mirror those considered in this paper, every committee member prefers *unanimity rule* to other k-majority rules—mechanisms which enact the policy if at least k members vote "yes". This is because, under unanimity, the most-demanding member[17] becomes the *de facto* dictator: others never vote "no" when she votes "yes", and when she votes "no", the policy is blocked. To persuade her, the lobbyist provides the committee more information than he would under any other k-majority rule. As such, the committee makes the best decisions under unanimity.

Proposition 2.2 establishes an analogous result for our setting.

Proposition 2.2. Let ε and $\overset{\bigstar}{\varepsilon}$ be two Blackwell experiments inducing the distributions μ and $\overset{\bigstar}{\mu}$ over interim beliefs, where supp $\overset{\bigstar}{\mu} = \{0, \overset{\bigstar}{q}\}$. If ε is Blackwell more informative than $\overset{\bigstar}{\varepsilon}$, the lobbyist expects a lower equilibrium payoff with μ than with $\overset{\bigstar}{\mu}$ under any voting mechanism.

[16] *The politician* in the former, *the Sender* in the latter.

[17] In their setting, committee members' upper thresholds of persuasion coincide with their thresholds of indifference, since they cannot acquire any private information. Then, the most-demanding member is simply the one who is least biased towards the policy; i.e., with the lowest u_i.

Proof. By Blackwell's Theorem (Blackwell (1953)),[18] if ε is Blackwell more informative than $\overset{\bigstar}{\varepsilon}$, then for any convex function ϕ over $\mathbb{R}$:

$$\int \phi(r)d\mu(r) \geq \int \phi(r)d\overset{\bigstar}{\mu}(r)$$

By Theorem 8 in Strassen (1965), this condition implies the existence of a belief martingale $\{\xi_1, \xi_2\}$ such that $\xi_1 \sim \overset{\bigstar}{\mu}$ and $\xi_2 \sim \mu$. Since $\{\xi_1, \xi_2\}$ is a belief martingale, $\mathbb{P}(\xi_2 = 0 \mid \xi_1 = 0) = 1$. Then, $\xi_2 \neq \xi_1$ only when $\xi_1 = \overset{\bigstar}{q}$. But by Claim 2.1, the interim belief $\overset{\bigstar}{q}$ leads to the enactment of the policy under any voting mechanism. So, the interim belief martingale ξ_2 can never increase the lobbyist's expected payoff above that induced by ξ_1.

□

Proposition 2.2, however, only provides a partial intuition for Theorem 2.1. Even though other mechanisms extract less information from the lobbyist than the $\overset{\bigstar}{m}$-dictatorship does, they may incentivise committee members to acquire more information. By doing so, they may help the committee accumulate more information overall, leading to better decisions.

What drives Theorem 2.1 is that the lobbyist *always* prevents this. He can always do so with the same experiment he would perform for a $\overset{\bigstar}{m}$-dictatorship. That experiment discourages committee members from acquiring further information and persuades them to enact the policy as frequently as the $\overset{\bigstar}{m}$-dictatorship would, regardless of the voting mechanism they use. He always wishes to do so, too. A $\overset{\bigstar}{m}$-dictatorship never blocks a good policy; so, the committee can improve its decisions only by blocking bad policies more often. So, the lobbyist never wishes the committee to take better decisions than the $\overset{\bigstar}{m}$-dictatorship would.

[18] The reader may note that Blackwell (1953) states the equivalence only for the standard measure of an experiment; i.e., the distribution of interim beliefs under the prior $p = 1/2$, here. However, it is easily confirmed that Blackwell (1953)'s proof for the "only if" direction holds for any the interim belief distribution generated by *any* prior $p \in [0, 1]$.

2.5 Discussion

In this Section, I conclude by discussing the robustness of Theorem 2.1 to various modelling assumptions.

Simplifying Assumptions

The model assumes that committee members' experiments are mutually independent conditional on the state. This means that committee members cannot directly gather information on their peers' private signals; they can only do so indirectly by gathering information about the unknown state that generates those signals. It also assumes that the lobbyist can only provide *public* information—send the same information to all committee members. It does not explicitly cover the case of *private* information provision, where the lobbyist may be allowed to communicate different bits of information to each committee member.

While both restrictions streamline exposition, neither affects the results presented in the paper. Clearly, neither restriction is relevant for how dictatorships behave; hence, for Lemma 2.1. Theorem 2.1 and Proposition 2.2 rest on the key insight that the information that persuades the $\overset{\bullet}{m}$-dictatorship persuades any other committee, too: it discourages any such committee from acquiring further information, and ensures that it enacts the policy at least as often as the $\overset{\bullet}{m}$-dictatorship. This insight hinges only on committee members' cost of acquiring information about the unknown state ω and their preferences about the policy; it is not sensitive to the removal of the aforementioned restrictions on the model. Thus, the results in this paper can easily be extended to a setting where committee members can acquire information about each others' signals, and the lobbyist can communicate privately with each member.

Equilibrium Selection

I focus on the lobbyist-preferred subgame perfect equilibria that different voting mechanisms induce. Under some voting rules—for instance, simple majority—this may lead to odd voting strategies. Particularly, committee members may vote to enact the policy regardless of their beliefs and preferences, simply because each believes that her vote will not be pivotal. To avoid such strategies, one may refine the set of equilibria further; for

instance, by excluding equilibria in weakly dominated voting strategies (e.g., like Alonso and Câmara (2016) and Heese and Lauermann (forthcoming)).

Such refinements on voting strategies have no bearing on the results presented here. They do not affect the equilibria dictatorships induce: dictators always vote for their preferred outcome. They do not affect the key insight behind Theorem 2.1 or Proposition 2.2 either: at the interim belief $\overset{\bigstar}{q}$ that persuades $\overset{\bigstar}{m}$ to enact the policy, any committee member prefers to see the policy enacted rather than blocked. So, information that persuades the $\overset{\bigstar}{m}$-dictatorship continues to persuade any other committee.

Nonetheless, the information strategies my equilibrium concept selects do play a key role in establishing Theorem 2.1 and Proposition 2.2. Following the interim belief $\overset{\bigstar}{q}$, no committee member acquires additional information—she instead votes to enact the policy. This information strategy is a best response given her belief that, in equilibrium, her peers will not acquire any private information either. Since these strategies secure the lobbyist's favourite outcome, they constitute the lobbyist-preferred subgame perfect equilibrium.

This, however, does not preclude subgame perfect equilibria where the committee *does* acquire private information following the interim belief $\overset{\bigstar}{q}$. Indeed, consider a committee member who instead believes that her peers will acquire additional private information at this interim belief. She no longer believes that the committee will decide based on the interim belief $\overset{\bigstar}{q}$; her peers may grow less certain that the policy was good, and vote against it. If they hold such private information, should she vote to allow her peers to block the policy? Or, should she counteract them, possibly reassured by her positive private information? To respond to this contingency, she may now find it useful to acquire private information herself.

This incentive to acquire information in order to best respond to others' renewed uncertainty mirrors the *encouragement effect*, first explored by Bolton and Harris (1999). Unlike their setting, committee members must *simultaneously* decide whether to acquire information, *before* they discover what their peers know. Nonetheless, an analogue of the encouragement effect may still emerge, thanks to members' incentives to best respond to their peers' informed votes.

My equilibrium selection rule precludes such equilibria—the lobbyist is best off when committee members vote in favour of the policy without acquiring further information.

Studying when these equilibria emerge and the outcomes they lead to is itself a very interesting and promising line of inquiry, left to future research.

The Lobbyist's Information Production Costs

The model assumes that the lobbyist can produce information freely—he can *costlessly* choose *any* Blackwell experiment about the state. This assumption has become standard and widespread in the literature following Kamenica and Gentzkow (2011)'s seminal paper. Nonetheless, it is a very powerful assumption, and a key driver of the results presented in this paper. To illustrate how it shapes the results, I extend the framework this paper analysed by constraining the lobbyist's ability to produce information for the committee.

Let E be a finite set of Blackwell experiments about the unknown state ω, and $\Delta_E \subset \Delta_p([0,1])$ be the set of interim belief distributions some experiment in E may generate. Suppose that the lobbyist's information strategy μ_l must now be an element of this set Δ_E, rather than any Bayes plausible distribution $\mu \in \Delta_p([0,1])$. Otherwise, assume that the setting remains the same.

Proposition 2.3. Suppose that $\overset{\star}{\mu} \in \Delta_E$, where supp $\overset{\star}{\mu} = \{0, \overset{\star}{q}\}$. Then, Lemma 2.1, Theorem 2.1 and Proposition 2.2 remain valid.

Proof. Proposition 2.2 trivially remains valid, as it does not depend on the set of experiments available to the lobbyist.

Lemma 2.1 establishes that under the $\overset{\star}{m}$-dictatorship, no strategy $\mu_l \in \Delta_p([0,1])$ delivers a greater equilibrium expected payoff than $\overset{\star}{\mu}$ for the lobbyist. Therefore, $\overset{\star}{\mu}$ remains optimal for the lobbyist when he is restricted to strategies $\mu \in \Delta_E$. As such, Lemma 2.1's characterisation of the equilibrium under the $\overset{\star}{m}$-dictatorship remains valid. Claim 2.1 and the ensuing argument used to establish Theorem 2.1 only rely on Lemma 2.1's characterisation of the equilibrium under a $\overset{\star}{m}$-dictatorship, the lobbyist's ability to employ the strategy $\overset{\star}{\mu}$, his preferences regarding policy outcomes, and committee members' best responses. Therefore, Theorem 2.1 also remains valid.

□

When $\overset{\star}{\mu} \notin E$, however, Proposition 2.3 fails. I show this through a simple counter-example.

Suppose there are two committee members, m_1 and m_2, who share the prior belief $p = 0.5$ with the lobbyist. m_1 and m_2 receive the payoffs $u_1 = 1/4$ and $u_2 = 2/3$ when a good policy is enacted. Thus, their thresholds of indifference are $(1 + u_1)^{-1} = 0.8$ and $(1 + u_2)^{-1} = 0.6$, respectively. Define $\pi_i(q)$ to be the payoff member m_i expects when her preferred decision is taken at posterior belief q:

$$\pi_1(q) := \begin{cases} 0 & q < 0.8 \\ 5/4 \times (q - 0.8) & q \geq 0.8 \end{cases} \qquad \pi_2(q) := \begin{cases} 0 & q < 0.6 \\ 5/3 \times (q - 0.6) & q \geq 0.6 \end{cases}$$

Let m_i's cost of acquiring information be defined by the convex function $c_i(q) := q^2 + \pi_i(q)$.

In turn, suppose that the lobbyist's information strategy must be an element of $\Delta_E = \{\mu_1, \mu_2\}$, where μ_1 and μ_2 are the unique distributions such that $\mu_1, \mu_2 \in \Delta_{0.5}([0, 1])$, supp $\mu_1 = \{0.2, 0.8\}$, and supp $\mu_2 = \{0, 0.6\}$.

In this counter example, m_1 is the most-demanding member of the committee: her upper threshold of persuasion equals $\bar{q}_1^* = 0.8$ against m_2's $\bar{q}_2^* = 0.6$. Yet, the m_1-dictatorship does not dominate the m_2-dictatorship.

To see this, first note that neither member finds it optimal to acquire any information under her dictatorship; both members have the equilibrium information strategy $\{\mu_i^*(\cdot \mid r) : \mu_i^*(\cdot \mid r) = \delta_r \text{ for all } r \in [0, 1]\}$. This is because m_i's equilibrium information strategy must solve the following problem:

$$\mu_i^*(\cdot \mid r) \in \underset{\mu \in \Delta_r([0,1])}{\arg\max} \int_{[0,1]} [\pi_i(q) - c_i(q)] \, d\mu(q) - [\pi_i(r) - c_i(r)]$$

The function $\pi_i(q) - c_i(q) = -q^2$ is strictly concave throughout its domain. Due to Jensen's Inequality, this implies that $\mu_i(\cdot \mid r) = \delta_r$ is the unique solution to this problem.

Thus, both members' upper thresholds of persuasion coincide with their thresholds of indifference: $\bar{q}_i^* = (1 + u_i)^{-1}$. Under the m_1-dictatorship, the strategy μ_2 never allows the lobbyist to enact the policy—its support lies strictly below m_1's upper threshold of indifference. In contrast, the strategy μ_1 enacts the policy when the interim belief $r = 0.8$ is realised. Hence, μ_1 is the lobbyist's equilibrium strategy under m_1's dictatorship. This

dictatorship achieves the following conditional probabilities:

$$\mathbb{P}_{\mu_1}(\text{enact} \mid \omega = \omega_1) = \mathbb{P}_{\mu_1}(r = 0.8 \mid \omega = \omega_1) = 4/5$$
$$\mathbb{P}_{\mu_1}(\text{block} \mid \omega = \omega_0) = \mathbb{P}_{\mu_1}(r = 0.2 \mid \omega = \omega_0) = 4/5$$

In contrast, μ_2 is the optimal strategy under the m_2-dictatorship. This strategy enacts the policy when the interim belief $r = 0.6$ is realised. This occurs with probability $5/6$, as one may easily verify using Bayes Rule. In contrast, μ_1 enacts the policy only when the interim belief $r = 0.8$ is realised. This occurs with probability $1/2$. Hence, μ_2 is the lobbyist's equilibrium strategy under m_2's dictatorship. This dictatorship achieves the following conditional probabilities:

$$\mathbb{P}_{\mu_2}(\text{enact} \mid \omega = \omega_1) = \mathbb{P}_{\mu_2}(r = 0.6 \mid \omega = \omega_1) = 1$$
$$\mathbb{P}_{\mu_2}(\text{block} \mid \omega = \omega_0) = \mathbb{P}_{\mu_2}(r = 0 \mid \omega = \omega_0) = 1/3$$

This verifies that m_1's dictatorship does not dominate m_2's dictatorship, and vice versa. Under m_2's dictatorship, the lobbyist faces a lower upper threshold of persuasion than under m_1's dictatorship. As such, he tailors the information he provides to m_2 so that she is likelier to enact the policy than m_1 is. Particularly, m_2 becomes likelier to enact a bad policy than m_1 is. However, she also becomes likelier to enact a good policy. Such a trade-off never emerges if the lobbyist can generate *any* interim belief distribution $\mu \in \Delta_p([0,1])$. He then ensures that dictators always enact a good policy.

References

Alonso, Ricardo and Odilon Câmara (Nov. 2016). "Persuading Voters". In: *American Economic Review* 106.11, pp. 3590–3605. DOI: 10.1257/aer.20140737.

Austen-Smith, David and Jeffrey S. Banks (1996). "Information Aggregation, Rationality, and the Condorcet Jury Theorem". In: *The American Political Science Review* 90.1, pp. 34–45. ISSN: 00030554, 15375943. DOI: 10.2307/2082796.

Bardhi, Arjada and Yingni Guo (2018). "Modes of persuasion toward unanimous consent". In: *Theoretical Economics* 13.3, pp. 1111–1149. DOI: 10.3982/TE2834.

Blackwell, David (1953). "Equivalent Comparisons of Experiments". In: *The Annals of Mathematical Statistics* 24.2, pp. 265–272. ISSN: 00034851. DOI: 10.1214/aoms/1177729032.

Bolton, Patrick and Christopher Harris (1999). "Strategic Experimentation". In: *Econometrica* 67.2, pp. 349–374. ISSN: 00129682, 14680262. DOI: 10.1111/1468-0262.00022.

Breyer, Stephen (1982). *Regulation and Its Reform*. Harvard University Press. ISBN: 9780674753754. URL: http://www.jstor.org/stable/j.ctv1smjtgr.

Caplin, Andrew and Mark Dean (Aug. 2013). *Behavioral Implications of Rational Inattention with Shannon Entropy*. Tech. rep. Working Paper 19318. National Bureau of Economic Research. DOI: 10.3386/w19318.

Caplin, Andrew, Mark Dean, and John Leahy (Aug. 2017). *Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy*. NBER Working Papers 23652. National Bureau of Economic Research, Inc. URL: http://www.nber.org/papers/w23652.

— (2022). "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy". In: *Journal of Political Economy* 130.6, pp. 1676–1715. DOI: 10.1086/719276.

Chan, Jimmy, Seher Gupta, et al. (2019). "Pivotal persuasion". In: *Journal of Economic Theory* 180, pp. 178–202. ISSN: 0022-0531. DOI: 10.1016/j.jet.2018.12.008.

Chan, Jimmy, Alessandro Lizzeri, et al. (2018). "Deliberating Collective Decisions". In: *The Review of Economic Studies* 85.2 (303), pp. 929–963. ISSN: 00346527, 1467937X. DOI: 10.1093/restud/rdx028.

Coglianese, Cary, Richard Zeckhauser, and Edward A. Parson (2004). "Seeking Truth for Power: Informational Strategy and Regulatory Policymaking". In: *Minnesota Law Review* 89.2, pp. 277–341.

Condorcet, Jean-Antoine-Nicolas de Caritat (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix.* Par.

Denti, Tommaso (Oct. 2022). "Posterior Separable Cost of Information". In: *American Economic Review* 112.10, pp. 3215–59. DOI: 10.1257/aer.20211252.

Feddersen, Timothy J. and Wolfgang Pesendorfer (1996). "The Swing Voter's Curse". In: *The American Economic Review* 86.3, pp. 408–424. ISSN: 00028282. URL: http://www.jstor.org/stable/2118204.

— (1997). "Voting Behavior and Information Aggregation in Elections With Private Information". In: *Econometrica* 65.5, pp. 1029–1058. ISSN: 00129682, 14680262. DOI: 10.2307/2171878.

— (1998). "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting". In: *American Political Science Review* 92.1, pp. 23–35. DOI: 10.2307/2585926.

Frankland, Mark (Feb. 3, 2004). *Lord Bullock of Leafield.* URL: https://www.theguardian.com/news/2004/feb/03/guardianobituaries.obituaries (visited on 07/08/2025).

Gerardi, Dino and Leeat Yariv (2008). "Information acquisition in committees". In: *Games and Economic Behavior* 62.2, pp. 436–459. ISSN: 0899-8256. DOI: 10.1016/j.geb.2007.06.007.

Gersbach, Hans (1995). "Information efficiency and majority decisions". In: *Social Choice and Welfare* 12.4, pp. 363–370. ISSN: 01761714, 1432217X. DOI: 10.1007/BF00186280.

Gershkov, Alex and Balázs Szentes (2009). "Optimal voting schemes with costly information acquisition". In: *Journal of Economic Theory* 144.1, pp. 36–68. ISSN: 0022-0531. DOI: 10.1016/j.jet.2008.02.004.

Heese, Carl and Stephan Lauermann (forthcoming). "Persuasion and Information Aggregation in Elections". In: *Journal of Political Economy*, null. DOI: 10.1086/736767.

Kamenica, Emir and Matthew Gentzkow (Oct. 2011). "Bayesian Persuasion". In: *American Economic Review* 101.6, pp. 2590–2615. DOI: 10.1257/aer.101.6.2590.

Li, Hao (2001). "A Theory of Conservatism". In: *Journal of Political Economy* 109.3, pp. 617–636. ISSN: 00223808, 1537534X. DOI: 10.1086/321018.

Matysková, Ludmila and Alfonso Montes (2023). "Bayesian persuasion with costly information acquisition". In: *Journal of Economic Theory* 211.C. DOI: 10.1016/j.jet.2023.105678.

Persico, Nicola (2004). "Committee Design with Endogenous Information". In: *The Review of Economic Studies* 71.1, pp. 165–191. ISSN: 00346527, 1467937X. DOI: 10.1111/0034-6527.00280.

Smorodinsky, Rann and Moshe Tennenholtz (2004). "Sequential information elicitation in multi-agent systems". In: *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*. UAI '04. Banff, Canada: AUAI Press, pp. 528–535. ISBN: 0974903906.

— (2006). "Overcoming free riding in multi-party computations—The anonymous case". In: *Games and Economic Behavior* 55.2. Mini Special Issue: Electronic Market Design, pp. 385–406. ISSN: 0899-8256. DOI: 10.1016/j.geb.2005.05.001.

Stephenson, Matthew C. (Apr. 2011). "Information Acquisition and Institutional Design". In: *Harvard Law Review* 124.6, pp. 1422–1483.

Strassen, V. (1965). "The Existence of Probability Measures with Given Marginals". In: *The Annals of Mathematical Statistics* 36.2, pp. 423–439. ISSN: 00034851. DOI: 10.1214/aoms/1177700153.

Chapter 3

Collusive College Admissions

Abstract

Two colleges compete for students by making costly investments. Students face uncertainty about colleges' quotas. I compare a regime where students can freely apply to both colleges with another where they can only apply to one. The former regime always generates a stable matching of students to colleges. However, colleges might be better off in the latter regime: the low-ranked college benefits from an *allocative effect* whereby it enrols better students, and the high-ranked college benefits from a *competitive effect* which decreases the competition for the best students. This incentivises colleges to coordinate their application policies and prevent students from applying to both. This finding helps explain the global persistence of college admission systems that generate unstable matchings.

*An earlier version of this paper was submitted to the University of Oxford in June 2020 for the degree of Master of Philosophy in Economics.

3.1 Introduction

Gale and Shapley (1962) showed that a stable matching exists in any college admissions problem, and proposed a mechanism which always implements it. Nonetheless, some college admission systems remain puzzlingly far from the stable mechanism they proposed. In the UK, undergraduate students are barred from applying both to Oxford and Cambridge in the same year.[1] Until 1994, South Korean students were limited to at most two applications, and could effectively apply to only one selective college since all held their admission exams on the same day (Avery, Lee, and Roth (2014)). Today, South Korean students remain limited to six applications (Eom (2024)).[2]

Two frictions—costly or restricted applications, and admission uncertainties—prevail in these admission systems. These frictions have drawn substantial attention in the Economics literature. The literature agrees on one prediction: when two colleges—one commonly preferred to ("ranked above") the other—compete for students, these two frictions may leave the low-ranked college better off at its competitor's expense. This occurs because these frictions force students to a trade-off: applying to a high-ranked college is more lucrative, but the low-ranked college presents a safer option. This may help the low-ranked college attract desirable students who would have been admitted by the high-ranked college, but chose not to apply there for fear of rejection.

This prediction emerges in various contexts. Avery and Levin (2010) show that when students can early-apply to a single college, admission uncertainties can generate a *strategic effect* which allows a lower-ranked college to attract desirable applicants away from its competitor. Chen and Kao (2014) show that with noisy screening, limited applications may help a low-ranked college get the best student through a coordination failure between students. Avery, Lee, and Roth (2014) similarly find that when students are unsure of their calibre, choosing an application date that conflicts with a higher-ranked competitor may help a low-ranked college attract top students. Chen and Kao (2023) empirically document that departments in a second-rank college in Taiwan indeed did strategically coincide their application dates with the best-regarded college, and that this improved

[1] See University of Oxford (2025).

[2] 4-year colleges founded for special purposes such as "industrial, open, online, military, science and technology, and music and art" are not bound by this limit (Eom (2024)).

their student intakes.

However, these studies treat students' preferences towards colleges—or, "college rankings"—as exogenous. In reality, colleges spend substantial effort towards improving or maintaining their rankings, which are often critical for students' application decisions. Indeed, due to their influence on students' decisions, college rankings have been accused of fuelling prestige wars, increasing education costs, and incentivising colleges to undertake unproductive—sometimes even illicit—endeavours to improve their standing.[3]

I study how these two frictions influence college admissions in a model where students' preferences between colleges are *endogenous*: colleges increase their appeal to students through costly investments. All students rank the college that invests more above its competitor. Students differ in their exogenous abilities, and both colleges prefer a higher-ability student to a lower-ability one. There is aggregate uncertainty about colleges' capacities; there are some applicants the college would admit when it has high, but not when it has low capacity. I compare two application regimes. Under the *Sequential Applications* regime, students apply to both colleges since colleges receive applications on different dates. Under the *Simultaneous Applications* regime, they only apply to one since colleges receive applications on the same date.

Echoing the literature, I find that the Simultaneous Applications regime may leave the low-ranked college better off due to an *allocative effect*: this college becomes a safe option for high-ability students who fear that the high-ranked college may not have enough capacity to admit them. Crucially, however, I find that the allocative effect also generates a *competitive effect* which benefits the high-ranked college, instead of harming it. Since the low-ranked college attracts some high-ability students as a safe option, it becomes less willing to compete with the high-ranked college for the best students. The high-ranked college concedes some high-ability students to its competitor, but faces significantly less competition for the rest. This competitive benefit always outweighs its allocative loss.

This dual benefit colleges reap comes at the expense of allocative fairness. Sequential Applications always generate a stable matching of students to colleges—justified envy never emerges as no student secures a worse outcome than a lower-ability peer. Under Simultaneous Applications, however, justified envy emerges: some high-ability students

[3]See Barkhorn (2014), Tierney (2013), and Trachtenberg (2011).

insure themselves by applying to the low-ranked college, but consequently, their lower-ability peers may enrol at the high-ranked college and secure better outcomes than them.

The best students are immune to this distortion: they always apply to, and enrol at, the high-ranked college. However, they are hurt by the competitive effect: the high-ranked college invests much less when it faces less competition from the low-ranked college for top students. Consequently, the best students enrol at the same college, but enjoy a much lower level of investment.

Motivated by this result, I then study the equilibrium of a game where colleges set their own application dates rather than having them exogenously determined. I find that when colleges are sufficiently likely to have low capacities, they always choose conflicting admissions dates and play the Simultaneous Admissions subgame. This is because when colleges are sufficiently likely to have low capacities, they are discouraged from competing for a large cohort of good applicants, and instead prefer to take advantage of the allocative and competitive effects. This resonates with and reinforces previous theoretical and empirical findings: *all* colleges—not just lower-ranked ones—might prefer to strategically impede students from applying widely.

Previous work on college admissions assumed students' preferences towards colleges to be exogenous. Therefore, the main result here was largely overlooked. Endogenising students' preferences towards colleges completes a missing piece of the puzzle. A consensus between colleges towards centralized, stable mechanisms might be impeded not only because low-ranked colleges might gain from instability, but also because this might help attenuate the competition to attract desired applicants.

Related Literature

College admissions are widely studied in Economics. Here, I review some of the work that relates most closely to this paper.

Avery and Levin (2010) study early admission policies using a model with two colleges and a continuum of students. Both colleges and students have private information: colleges can screen their applicants' abilities and applicants know their fit for colleges. Both are payoff relevant for colleges. They find that early admissions without enrolment commitments generate a *sorting effect* which benefits both colleges: students credibly signal

their fit for a college through their early application choices. However, with enrolment commitments, a *strategic effect* also emerges. This helps the low-ranked college benefit at its competitor's expense by capturing cautious but high-ability applicants in the early applications stage. The existence and strength of these effects are determined by students' exogenous preferences for colleges.

Chen and Kao (2014) illustrate this strategic effect in a different context. They study an environment with a single application limit and noisy screening which causes enrolment uncertainty. They show that a coordination problem can cause the better student to be matched to the low-ranked college, as close ranked students avoid competing in applications. Again, this benefits the low-ranked college at the expense of the high-ranked college.

Avery, Lee, and Roth (2014) study a 1994 policy change in South Korea. In their model, students are uncertain either about their common or idiosyncratic values for colleges. Colleges set admission dates sequentially, which determines whether students are restricted to a single application. They find that when common values are unknown to students, a lower ranked college can benefit from date conflicts by capturing safety applicants revealed to have high ability. Key to whether the lower ranked college benefits from timing conflicts is whether the idiosyncratic or common values are unknown, and how close students regard the two colleges. Their focus is a policy change affecting the former, but they take the latter as given.

In a later paper, Chen and Kao (2023) study an educational reform in Taiwan which allowed departments in a second-ranked college to set their application dates freely. Their theoretical model predicts that a second-ranked college can attract applicants away from its competitor by preventing students from applying to both. Their empirical findings supports this prediction: they find that some departments moved their application dates to coincide with that of the best-ranked college, and that this led to an improvement in their student intake. Their theoretical model predicts that this benefits the second-ranked college at the expense of the best-ranked college. The predictions of my model align with their findings. However, my model predicts that such strategic distortions might benefit *both* colleges, and have worse consequences for students.

In less direct relation, Chade, Lewis, and Smith (2014) analyse college admissions in

a demand and supply framework, focusing on students' strategic application behaviour. They incorporate the frictions analysed here: applications are costly, and outcomes are uncertain as screening is noisy. They find that "non-sorting" can occur in equilibrium: when the lower ranked college compares closely to the top ranked, weaker students might apply more aggressively than their stronger peers. How colleges compare is again exogenous.

This literature focuses on the same frictions this paper does. However, this paper differs in the way it models the uncertainty students face regarding their enrolment outcomes: rather than noisy screening of students' abilities, this paper assumes that enrolment uncertainty emerges from aggregate uncertainty about colleges' capacities. This makes the model significantly simpler and more tractable.

These studies show that decentralized admissions with frictions can benefit low-ranked colleges by helping them enrol unachievable students. Whether this happens, and its extent, depends critically on how strong students' preferences between colleges are. My main contribution to this literature is to endogenise colleges' competition for students. In my model, alongside their admission timings, colleges choose how much costly investment to undertake, determining students' preferences towards them. This proves useful not only to better understand the nature of competition between colleges for students, but also highlights an overlooked source of gains for colleges from frictions. High-ranked colleges need not necessarily lose from frictions when their rival captures some coveted students. As this attenuates competition for the best students, both colleges can gain in equilibrium.

This paper also relates to a strand of literature exploring matching markets with ex-ante investments. Felli and Roberts (2016) analyse a one-to-one matching market where heterogenous buyers and sellers make ex-ante investments and are matched by Bertrand competition. Agents' qualities are determined by their investments and exogenous types. Both sides' qualities are complementary for the surplus the match generates, and critically, utilities are transferable. They focus on the ways efficiency can be impeded; a coordination problem can prevent assortative matching, and a hold-up problem can cause underinvestment in equilibrium.

Burdett and Coles (2001) study a similar but dynamic marriage market with heterogenous agents and non-transferable utilities. Agents can invest ex-ante to improve their *pizazz* in hopes for a better match, but otherwise irrelevant to their payoff. They find an

over-investment equilibrium exists as agents do not account that their self-improvement, making them more selective, makes those who do not improve in the opposite side of the market worse off.

In my model, utility is non-transferable between colleges and students, unlike Felli and Roberts (2016). As in Burdett and Coles (2001), colleges' investments are their bids for potential matches. However, rather than on potential inefficiencies arising in a competitive market, I focus on how a change in the matching mechanism changes colleges' incentives to invest, and the subsequent allocations and payoffs, in a one-to-many matching market.

Outline

The rest of the paper is structured as follows: Section 3.2 presents the model. Section 3.3 analyses the Sequential Applications game, and Section 3.4 analyses the Simultaneous Applications game. Section 3.5 studies the equilibrium of the college admissions game where colleges choose their admission dates. Section 3.6 concludes. Some results not central to the discussion are omitted in the main text. They are reserved for Section 3.7. In many instances, proofs for the results presented in the main text are reserved for Section 3.8.

3.2 The Model

There are two ex-ante identical colleges A and B and a unit mass of students with types $a \sim U[0,1]$, where a represents their intrinsic ability. Constrained by their ex-ante unknown capacities κ, colleges enrol students to produce human capital by combining their costly investment with their students' abilities. Ability and investment are complementary inputs in human capital production.[4] The total human capital a college produces is given by:

$$k(\xi, F) := \xi \times \int_0^1 a \times dF(a)$$

where ξ is the costly investment of the college and $F : \mathscr{B}([0,1]) \to [0,1]$ is the measure of its students' abilities. This college's payoff is the human capital it produces less its

[4]The marginal product of investment increases in ability, and vice versa. Where $k(a,h)$ is the human capital production function, $\frac{\partial^2 k(a,\xi)}{\partial a \partial \xi} > 0$.

investment cost:

$$\Pi(\xi) = k(\xi, F) - \xi^2$$

.

There are no complementarities between students; i.e., colleges' preferences over sets of students are responsive, and a higher-ability student is commonly preferred to a low-ability student. I therefore refer to higher-ability (lower-ability) students as "better" ("worse") students. Students' payoff is the human capital they help produce. A student with ability a and matched to a college with investment ξ receives the payoff:

$$u(a, \xi) := \xi \times a$$

A student matched to no college gets a payoff of 0. Therefore, all students prefer the college that offers higher investment, and do not regard any college as unacceptable. Hence, I refer to the college offering higher investment as the "high-ranked" college, and to its competitor as the "low-ranked" college.

The surplus from human capital production cannot be redistributed with side payments; i.e., utilities are non-transferable. All players are assumed to be risk-neutral, and each maximises their expected payoff.

Students are matched to colleges through a sequential move, symmetric, and perfect information college admissions game. The game proceeds with the following timing:

1. A chooses and announces a date $t_A \in \{1, 2\}$ to receive applications.
2. B chooses and announces a date $t_B \in \{1, 2\}$ to receive applications.
3. A chooses and announces its investment $\alpha \geq 0$. It pays the sunk cost α^2.
4. B chooses and announces its investment $\beta \geq 0$. It pays the sunk cost β^2.
5. Students choose which college to apply to at dates 1 and 2. They can apply to a single college at either date.
6. Nature chooses the capacity κ for both colleges. $\mathbb{P}(\kappa = \kappa_L) = p \in (0, 1)$, and $\mathbb{P}(\kappa = \kappa_H) = 1 - p \in (0, 1)$, where $\frac{1}{2} \geq \kappa_H > \kappa_L$.

7. The Gale and Shapley (1962) Deferred Acceptance algorithm, as formalised by Abdulkadiroğlu, Che, and Yasuda (2015) for settings with a continuum of students, is run between the applicants and colleges to match them. Any student who has not applied to a college is treated as having declared it unacceptable.

I use Subgame Perfect Nash Equilibria (SPNE) as the solution concept. In Sections 3.3 and 3.4, I study the SPNE of the subgames with $t_A \neq t_B$ and $t_A = t_B$, respectively. I refer to the former as the *Sequential Applications* game, and the latter as the *Simultaneous Applications* game. In Section 3.5, I study the SPNE of the whole college admissions game. I assume that, when available, players play their weakly dominant strategies in equilibrium. I also assume that, when a college expects equal payoffs from two strategies, it chooses the strategy with the higher investment level.

3.3 Sequential Applications

Consider the *Sequential Applications* subgame, induced when $t_A \neq t_B$. Since applications are held at different dates, students' weakly dominant strategy is to apply to both colleges. When students have strict preferences between colleges ($\alpha \neq \beta$), the positive assortative matching of students to colleges is the unique stable matching of the college admissions game.

Definition 3.1. Let $\alpha > \beta$. The positive assortative matching of students to colleges is described by $\mu^s : [0,1] \times \{\kappa_L, \kappa_H\} \to \{\emptyset, A, B\}$, which maps a student's ability and realised capacities to an assignment for the student:

$$\mu^s(a;\kappa) := \begin{cases} A & a \geq 1-\kappa \\ B & a \in [1-2\kappa, 1-\kappa) \\ \emptyset & a < 1-2\kappa \end{cases}$$

where unmatched students are assigned to $\emptyset$. When $\beta > \alpha$, students in A are swapped with those in B.

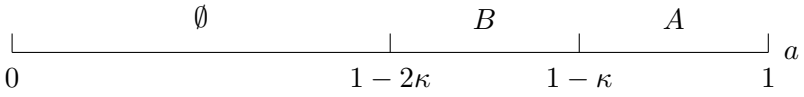

Figure 3.3.1: The Positive Assortative Matching $\mu^s(a, \kappa)$

Lemma 3.1. The positive assortative matching μ^s is the unique stable matching of the college admissions game when $\alpha \neq \beta$. Consequently, it is the unique equilibrium outcome of the Sequential Admissions game when $\alpha \neq \beta$.

As students apply to both colleges, the Gale and Shapley (1962) Deferred Acceptance algorithm implements the unique stable matching μ^s irrespective of the realized capacity. Without limited applications, capacity and enrolment uncertainty become irrelevant for students' decisions and the resulting matching—enrolment uncertainty cannot preclude stability on its own.

Continue to assume that A invests more than B; i.e., $\alpha > \beta$. Given $\alpha > \beta$, denote the set of students college i enrols when the realised capacities are κ as $(\mu^s)^{-1}(i, \kappa)$. Conditional on their investments and realised capacities, colleges A and B expect the payoffs $\Pi^s_A(\alpha \mid \kappa)$ and $\Pi^s_B(\beta \mid \kappa)$ when students are positive assortatively matched to colleges:

$$\Pi^s_A(\alpha, \kappa) := \alpha \times \int_{1-\kappa}^{1} a.da - \alpha^2 \qquad \Pi^s_B(\beta, \kappa) := \beta \times \int_{1-2\kappa}^{1-\kappa} a.da - \beta^2$$

Given their beliefs about what realised capacities will be, colleges A and B expect the payoffs $\Pi^s_A(\alpha) := p \times \Pi^s_A(\alpha \mid \kappa_L) + (1-p) \times \Pi^s_A(\alpha \mid \kappa_H)$ and $\Pi^s_B(\beta) := p \times \Pi^s_B(\beta \mid \kappa_L) + (1-p) \times \Pi^s_B(\beta \mid \kappa_H)$ given their investment levels $\alpha > \beta$. Figure 3.3.2 plots these investment curves as a function of colleges' investments.

The functions $\Pi^s_A(\alpha)$ and $\Pi^s_B(\beta)$ attain their respective global maxima $\bar{\Pi}^s_A$ and $\bar{\Pi}^s_B$ at the investment levels $\bar{\alpha}^s$ and $\bar{\beta}^s$, where these investment levels are given as:

$$\bar{\alpha}^s := \frac{p \times \int_{1-\kappa_L}^{1} a.da + (1-p) \times \int_{1-\kappa_H}^{1} a.da}{2} \qquad \bar{\beta}^s := \frac{p \times \int_{1-2\kappa_L}^{1-\kappa_L} a.da + (1-p) \times \int_{1-2\kappa_H}^{1-\kappa_H} a.da}{2}$$

The plot also depicts the highest investment α for which college A, which enrols better students than college B, expects a weakly higher payoff than $\bar{\Pi}^s_B$; i.e., the highest expected

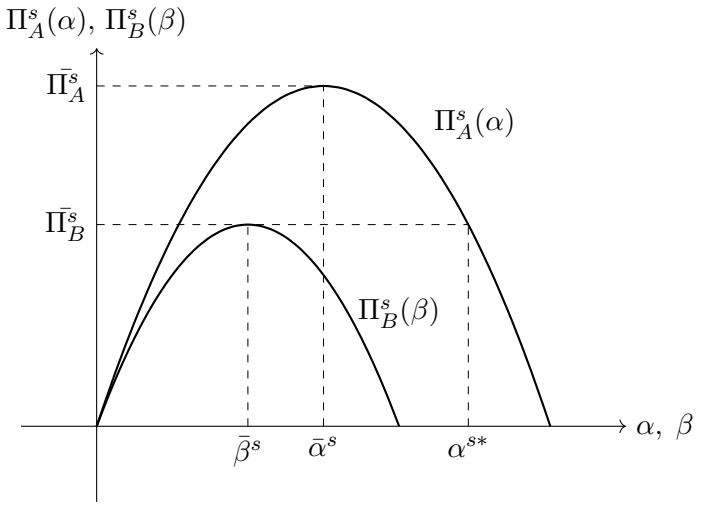

Figure 3.3.2: Colleges' Expected Payoffs under Sequential Applications

payoff B can secure given the students it expects to enrol. I denote this investment level as $\alpha^{s*} := \max\{\alpha : \Pi^s_A(\alpha) \geq \bar{\Pi}^s_B\}$. Proposition 3.2 then describes the equilibrium under the Sequential Applications game.

Proposition 3.2. In the unique equilibrium of the Sequential Applications game, college A invests α^{s*}, B invests $\bar{\beta}^s$, and the stable matching μ^s is realised.

College A faces two options. If it invests below α^{s*}, college B will best respond by posting a higher investment level and securing a payoff above $\bar{\Pi}^s_B$. A's expected payoff will then equal $\Pi^s_B(\alpha)$, since it will obtain the second-best group of students. The payoff function $\Pi^s_B(\cdot)$ is maximised at the investment level $\bar{\beta}^s < \alpha^{s*}$; so, the optimal investment level below α^{s*} is $\bar{\beta}^s$. This secures A an expected payoff of $\bar{\Pi}^s_B$.

Alternatively, A could invest weakly above α^{s*}. If college B posts a higher investment level, it will secure an expected payoff strictly below $\bar{\Pi}^s_B$. Therefore, B will best respond by setting its investment to $\bar{\beta}^s$, and securing a payoff of $\bar{\Pi}^s_B$. College A, attracting the best students, will then secure an expected payoff of $\Pi^s_A(\alpha)$. Since this function is strictly decreasing above α^{s*}, the optimal investment level which weakly exceeds α^{s*} is, again, α^{s*}. This secures college A an expected payoff of $\bar{\Pi}^s_B$. College A is then indifferent between the investment levels $\bar{\beta}^s$ and α^{s*}. As I focus on equilibria where colleges break ties in favour of higher investment levels, the latter investment level becomes A's equilibrium strategy.

3.4 Simultaneous Applications

3.4.1 Equilibrium Application Strategies and Matchings

Now consider the *Simultaneous Applications* subgame, induced when $t_A = t_B$. Applying to some college is still weakly dominant for students; but now, they must choose which college to apply to. I denote a student's application strategy—the probability that he applies to college A given his ability a—as $\sigma(a)$. An applicant to a college enrols there unless the mass of higher-ability applicants to that college meets or exceeds its capacity κ. Given this, let $\rho_i(a, \sigma)$ denote the probability that a student of ability a enrols to college i upon her application given her peers' application strategies σ. A student, then, strictly prefers to apply to college A instead of B if and only if $\rho_A(a, \sigma) \times \alpha > \rho_B(a, \sigma) \times \beta$.

The discussion that follows continues to assume that $\alpha > \beta$; the case $\beta > \alpha$ simply yields the symmetric analysis. In Sections 3.4.1.1 and 3.4.1.2, I analyse students' equilibrium application strategies and the resulting matching in two subgames: (i) $(1-p)\times\alpha > \beta$, and (ii) $(1-p)\times\alpha < \beta < \alpha$, respectively. I then uncover the conditions under which either subgame is played in equilibrium.

Students' strategies differ crucially between the subgames (i) and (ii). Consider a student whose ability places her in the top $\kappa_H{}^{\text{th}}$ quantile of students. She knows that the high-ranked college is certain to admit her when it has high capacity, but not when it has low capacity. In subgame (i), she finds the investment premium the high-rank college A offers sufficient to justify a risk of rejection, and optimally applies to A. I call such a student a *gambler*. In subgame (ii), the same student would prefer to apply to the low-rank college B: A's investment premium no longer offsets the safe alternative that B offers. I call such a student an *insurer*. As they incentivise different application behaviour, these two subgames induce different matchings of students to colleges, and consequently, different equilibrium payoffs for the high- and low-ranked colleges.

3.4.1.1 High Difference Between Investments: $(1-p) \times \alpha > \beta$

Let $(1-p) \times \alpha > \beta$; this can be thought as the scenario where the difference between colleges' investment levels is high. Clearly, students with types $a \geq (1-\kappa_L)$ optimally

apply to the high-rank college, A; colleges' capacities always weakly exceed κ_L, so, no college can ever fill its capacity with students preferable to these top applicants. As these students apply to A, students with types $(1-\kappa_L) \geq a \geq (1-\kappa_H)$ face a trade-off. Such students would be admitted by the high-rank college when $\kappa = \kappa_H$, but not when $\kappa = \kappa_L$. College B could offer them better admission prospects. Nonetheless, they are not deterred from applying to A; the investment premium A offers is sufficient to offset the risk of rejection. Thus, when the difference between colleges' investment levels is high, these students *gamble* on A. Since these students gamble on A, no other student has a chance to be admitted by A. Without loss, all such students apply to college B.

These equilibrium strategies give rise to a unique matching of students to colleges in equilibrium. Proposition 3.3—illustrated in Figure 3.4.1—summarizes it.

Proposition 3.3. Let the matching $\mu^{hi}(a,\kappa)$ be defined as:

$$\mu^{hi}(a,\kappa_L) := \begin{cases} A & a \geq 1-\kappa_L \\ \emptyset & a \in [1-\kappa_H, 1-\kappa_L) \\ B & a \in [1-\kappa_L-\kappa_H, 1-\kappa_H) \\ \emptyset & a < 1-\kappa_L-\kappa_H \end{cases}$$

$$\mu^{hi}(a,\kappa_H) := \begin{cases} A & a \geq 1-\kappa_H \\ B & a \in [1-2\kappa_H, 1-\kappa_H) \\ \emptyset & a < 1-2\kappa_H \end{cases}$$

When $(1-p)\times\alpha > \beta$, μ^{hi} is the unique matching that emerges in the equilibrium of the Simultaneous Admissions game. When $\kappa = \kappa_H$, this matching is pairwise stable; but when $\kappa = \kappa_L$, it exhibits justified envy.

When $\kappa = \kappa_H$, the matching μ^{hi} replicates the stable (positive assortative) matching μ^s. However, when $\kappa = \kappa_L$, students who gamble for the high-ranked college are left unmatched, even though worse students enrol to college B. This generates justified envy; so, when $\kappa = \kappa_L$, the matching μ^{hi} is not pairwise stable.

In this subgame, the high-ranked college enrols the same students it would in the

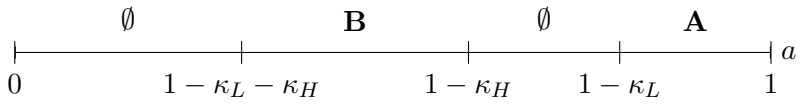

(a) $\mu^{hi}(a, \kappa_L)$

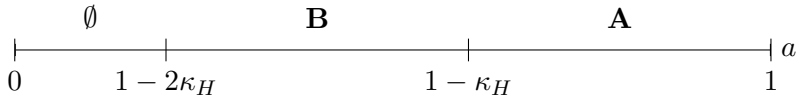

(b) $\mu^{hi}(a, \kappa_H)$

Figure 3.4.1: The Equilibrium Matching in Simultaneous Adm.: $(1-p) \times \alpha > \beta$

Sequential Applications game. Therefore, conditional on its investment level, its expected payoff is the same as it is under Sequential Applications. This comes at the expense of the low-ranked college's intake quality. In the Sequential Applications game, this college enrols the best students among those who are not admitted by the high-ranked college. In this subgame, it fails to enrol some such students, since those students apply exclusively to the high-ranked college. Since the low-ranked college expects to enrol worse students than it would in the Sequential Applications game, it also expects a lower payoff conditional on its investment.

3.4.1.2 Low Difference Between Investments: $\boldsymbol{\alpha > \beta > (1-p) \times \alpha}$

Let us now consider $\alpha > \beta > (1-p) \times \alpha$; this can be thought as the scenario where the difference between colleges' investment levels is low. The optimal application strategy for students with types $a \geq (1 - \kappa_L)$ is clearly identical. However, for those who would be admitted to A only when $\kappa = \kappa_H$, but are sure admits in B, the optimal strategy changes. The investment premium A offers no longer suffices to motivate these students to gamble for the high-ranked college.

Students with types $1 - 2\kappa_L \leq a < 1 - \kappa_L$ are sure admits in B. They are admitted to A with probability at most $(1-p)$: when $\kappa = \kappa_L$, A fills its capacity with students it favours over them. Therefore, their expected payoff from an application to B exceeds that from A; $(1-p) \times \alpha < \beta$. Hence, they *insure* themselves against admission uncertainty by

applying to B, even though they would be admitted to A with positive probability.

Given the optimal strategies of students with $a \geq 1-2\kappa_L$, those with types $1-\kappa_L-\kappa_H \leq a < 1-\kappa_H$ face equal admission chances at either college: either will admit them if and only if $\kappa = \kappa_H$. Hence, applying to A is clearly optimal for them. Remaining students have no admission chances at A; without loss, they apply to B.

These equilibrium strategies give rise to a unique matching of students to colleges in equilibrium. Proposition 3.4—illustrated in Figure 3.4.2—summarises it.

Proposition 3.4. Let the matching $\mu^{lo}(a,\kappa)$ be defined as:

$$\mu^{lo}(a,\kappa_L) := \begin{cases} A & a \geq 1-\kappa_L \\ B & a \in [1-2\kappa_L, 1-\kappa_L) \\ \emptyset & a < 1-2\kappa_L \end{cases}$$

$$\mu^{lo}(a,\kappa_H) := \begin{cases} A & a \in [1-\kappa_L, 1] \cup [1-\kappa_L-\kappa_H, 1-2\kappa_L) \\ B & a \in [1-2\kappa_L, 1-\kappa_L) \cup [1-2\kappa_H, 1-\kappa_L-\kappa_H) \\ \emptyset & a < 1-2\kappa_H \end{cases}$$

When $\alpha > \beta > (1-p)\times\alpha$, μ^{lo} is the unique matching that emerges in the equilibrium of the Simultaneous Admissions game. When $\kappa = \kappa_L$, this matching is pairwise stable; but when $\kappa = \kappa_H$, it exhibits justified envy.



(a) $\mu^{lo}(a,\kappa_L)$

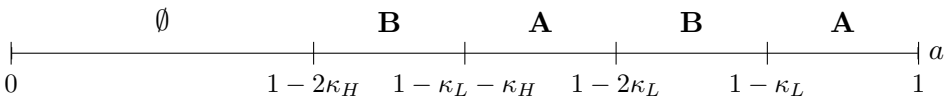

(b) $\mu^{lo}(a,\kappa_H)$

Figure 3.4.2: The Equilibrium Matching in Simultaneous Adm.: $\alpha > \beta > (1-p)\times\alpha$

Justified envy emerges here, too. Previous gamblers now *insure*; they apply to the low-ranked college, which presents a safer option. This guarantees that they are never left unmatched. When $\kappa = \kappa_L$, the equilibrium matching replicates $\mu^s(a, \kappa_L)$; hence, it is stable. However, since they only apply to the low-ranked college B, the high-ranked college A admits students worse than them when $\kappa = \kappa_H$. This causes justified envy.

In this subgame, the high-ranked college cannot attract every student it would like to enrol. The very best students, who face no admission uncertainty, still apply here. However, the best students among those who remain apply to the low-ranked college instead; they find the investment premium it offers too little to risk rejection when $\kappa = \kappa_L$. This improves the set of students the low-ranked college expects to enrol. It locks in the students its rival would like to admit when $\kappa = \kappa_H$ as insurance applicants. In expectation, these students have higher abilities than those it would admit in Sequential Applications. Therefore, it expects higher payoff conditional on its investment.

3.4.2 Investments and Payoffs in Simultaneous Applications

Unlike in Sequential Applications, the difference between the investments of the high- and low-ranked colleges influences the students these colleges are matched to in Simultaneous Applications. Section 3.4.1 analysed two subgames where the high and low-ranked colleges are matched to different sets of students, depending on how their investment levels compare.

Given the students the high- and low-ranked colleges secure in each subgame, we can derive the payoffs they expect given their investments. It is helpful to first introduce some notation. As I have done so far, I assume that A is the high-ranked college in the discussion and notation below—the symmetric discussion obtains when B is the high-ranked college instead. For each realised capacities κ, denote the set of students the high-ranked college obtains in subgame (i) as $(\mu^{hi})^{-1}(A, \kappa)$. For the low-ranked college, this is $(\mu^{hi})^{-1}(B, \kappa)$. For subgame (ii), use the analogous expressions, $(\mu^{lo})^{-1}(A, \kappa)$ and $(\mu^{lo})^{-1}(B, \kappa)$. Then, we can define $\Pi_A^{hi}(\alpha)$ and $\Pi_B^{hi}(\beta)$ to be the payoffs these colleges expect in subgame (i),

given their investment levels α and β:

$$\Pi_A^{hi}(\alpha) := \alpha \times \left[p \times \int\limits_{(\mu^{hi})^{-1}(A,\kappa_L)} a.da + (1-p) \times \int\limits_{(\mu^{hi})^{-1}(A,\kappa_H)} a.da \right] - \alpha^2$$

$$\Pi_B^{hi}(\beta) := \beta \times \left[p \times \int\limits_{(\mu^{hi})^{-1}(B,\kappa_L)} a.da + (1-p) \times \int\limits_{(\mu^{hi})^{-1}(B,\kappa_H)} a.da \right] - \beta^2$$

The analogous expressions, which I will denote as $\Pi_A^{lo}(\alpha)$ and $\Pi_B^{lo}(\beta)$, define these colleges' expected payoffs in subgame (ii). Each of these payoff curves is strictly concave, thus has a unique maximiser. Denote these unique maximisers as $\bar{\alpha^{hi}}$, $\bar{\beta^{hi}}$, and $\bar{\alpha^{lo}}$, $\bar{\beta^{lo}}$. These are the *optimal levels of investment* for each corresponding set of applicants. Similarly, denote the maximised values of these payoff curves as $\bar{\Pi_A^{hi}}$ and $\bar{\Pi_B^{hi}}$, and $\bar{\Pi_A^{lo}}$ and $\bar{\Pi_B^{lo}}$.

Now, consider what B's best response should be to any investment level α set by A. We can break this down to two questions:

1. Given A's investment α, which applicants should B aim to attract?

2. Which investment level β, among those that attract those applicants, maximise B's payoff?

There is an *optimal level of investment* for each set of applicants B may want to attract; given those applicants, this level of investment maximises its payoff. However, depending on what A's investment α is, B may not be able to attract those applicants at the corresponding optimal level of investment. When that level of investment is too low to attract the corresponding applicants, the best B can do is to set the lowest investment that would attract them; this is because the payoff curve for each set of applicants is strictly concave. To reflect this idea, I define *constrained expected payoff* functions for each of the four sets of applicants $(\mu^{hi})^{-1}(A,\kappa)$, $(\mu^{hi})^{-1}(B,\kappa)$, $(\mu^{lo})^{-1}(A,\kappa)$, and $(\mu^{lo})^{-1}(B,\kappa)$ that college B may wish to attract:

$$\widetilde{\Pi_A^{hi}}(\alpha) := \Pi_A^{hi}\left(\max\left\{\bar{\alpha^{hi}}, \frac{\alpha}{1-p}\right\}\right) \qquad \widetilde{\Pi_A^{lo}}(\alpha) := \Pi_A^{lo}\left(\max\left\{\bar{\alpha^{lo}}, \alpha\right\}\right)$$

$$\widetilde{\Pi_B^{hi}}(\alpha) := \Pi_B^{hi}\left(\max\left\{\bar{\beta^{hi}}, \alpha \times (1-p)\right\}\right) \qquad \widetilde{\Pi_B^{lo}}(\alpha) := \Pi_B^{lo}\left(\bar{\beta^{lo}}\right)$$

These functions map college A's investment level α to an upper bound for B's expected payoff for each corresponding set of applicants, given B must invest enough to attract that set of applicants.

In Section 3.7, I prove that the upper envelope of these four functions describes B's best response to any investment level α college A may set. As such, this upper envelope also describes the set of applicants, and therefore payoff, college A may expect in equilibrium, for each investment level α it may set.

The upper envelope of the four constrained expected payoff functions takes a different shape for every probability $p \in [0, 1]$ with which Nature may choose a capacity $\kappa = \kappa_L$. Furthermore, the analytical expressions for each of these four functions are complex; this makes it difficult to pin down and characterise the equilibrium for every possible parameter value. Nonetheless, Proposition 3.5 establishes that when p is sufficiently large, this equilibrium induces a low quality differential between colleges.

Proposition 3.5. There exists some $\bar{p} \in (0, 1)$ such that for all $p > \bar{p}$, the unique equilibrium of the Simultaneous Admissions game is such that:

- A and B's equilibrium investments α^* and β^* are such that $\alpha^* > \beta^* > (1 - p) \times \alpha^*$,
- the realised matching of students to colleges is $\mu^{lo}(a, \kappa)$,
- the expected payoff of both colleges is $\bar{\Pi}_B^{lo}$,
- relative to the equilibrium investments in the Sequential Admissions game, college B invests more, but college A invests less.

In this equilibrium, both colleges obtain a higher payoff relative to their payoffs in the Sequential Applications game. The low-ranked college, B, owes this to an *allocative benefit*. Under the Simultaneous Applications game, the best students among those who are not sure admits to the high-ranked college face a trade-off: either they gamble on the high-ranked college A and risk rejection when $\kappa = \kappa_L$, or they insure by applying to B and guarantee a spot at the low-ranked college. When colleges' investment levels are close, such students prefer to insure. College B locks these students' applications in: in the Sequential Applications game, A would enrol these students provided $\kappa = \kappa_H$. Here, B secures these

students regardless of the realised capacities. Since it gets better students, it invests more than it would in the Sequential Applications game.

The allocative benefit B obtains harms the quality of college A's intakes; when the capacities κ_H are realised, college A enrols worse students than the insurers who exclusively apply to B. However, B's allocative benefit generates a *competitive benefit* for college A. Under Sequential Applications, B had to outcompete A to attract the insurers. Here, it no longer has to do so. This reduces its incentives to outcompete A; doing so will still win the top ${\kappa_L}^{\text{th}}$ quantile of applicants over, but lose the insurers to its competitor. This allows A to deter B from competing for the top students at much less cost. As a result, it obtains a higher payoff than it did under Sequential Applications, but at much less cost.

High values of p facilitate this equilibrium. When p is low, colleges believe that capacities are significantly likely to realise as κ_H; this intensifies the incentive each college has to outcompete the other by a wide margin and capture all the best applicants.

3.5 Equilibrium when Colleges Set Application Dates

Proposition 3.5 establishes that for sufficiently high levels of p, both colleges secure higher payoffs in the Simultaneous Applications game than they do in the Sequential Applications game. Backward induction then immediately reveals that for such levels of p, only the Simultaneous Applications game is played in the equilibrium of the full game where colleges set their admission dates themselves.

Proposition 3.6. Suppose p is such that Proposition 3.5 holds. Then, in the equilibrium of the college admissions game, colleges set the same admission dates, $t_A^* = t_B^*$, and the Simultaneous Admissions game is played in the ensuing subgame.

In such equilibria, no college wishes a stable matching mechanism to operate. The Simultaneous Applications game generates student-college matchings with justified envy by placing some students coveted by the high-ranked college to college B. This benefits both colleges: the low-ranked college benefits from an allocative effect, and the high-ranked college benefits from the alleviated competition this enables. This comes at the expense of students who submit insurance applications: when capacities realise as κ_H, they are the ones who experience justified envy for not enrolling at the high-ranked college. The top

$\kappa_L{}^{\text{th}}$ quantile of students also suffer: they enrol at the high-ranked college in every game, but receive a lower level of investment in the Simultaneous Applications game.

3.6 Conclusion

Gale and Shapley 1962 show that a stable matching exists in any college admissions game, and proposed a simple mechanism that always implements it. Yet, many college admission systems are still puzzlingly far from this mechanism, exhibiting two key frictions: costly, or restricted applications, and admission uncertainties. Previous studies converged on the prediction that these frictions might benefit a low-ranked college by distorting stability in its favour. If students' preferences between this college and its high-ranked rival are not sufficiently strong, the low-ranked college can attract cautious applicants its rival would admit. This comes at the expense of the high-ranked college, which is deprived of these desirable applicants.

These studies treat students' preferences towards colleges— colleges' rankings—as exogenous. However, colleges spend substantial effort to improve their rankings and appeal to students by various means. These efforts are not limited to strategically setting their admission policies, which previous studies focus on. They also involve substantial productive or unproductive investments.

With this motivation, I extend this literature by endogenising students' preferences for colleges. I analyse a model where colleges compete for better applicants by investing in a human capital technology. Like earlier models, I find that when colleges are close competitors, the low-ranked college can benefit from limited applications and students uncertain about their admission prospects. These frictions distort the stable matching in its favour, helping it attract cautious applicants desirable by its rival. A direct implication is that the quality of the students the high-ranked college expects deteriorates.

However, endogenising students preferences yields the novel insight that this can improve the high-ranked college's payoff as well. This owes to a favourable competitive effect: as the low-ranked college gets better students, it becomes less willing to compete for the best ones. Consequently, the high-ranked college also gains from this ameliorated competition, despite getting worse students. Matching frictions can act as blessings cushioning

competition for both colleges.

I find that this outcome characterises the equilibrium when matching frictions are present if admission uncertainty is such that competing for all the best applicants is sufficiently unprofitable. In this equilibrium, both colleges are made better off relative to the equilibrium where no matching frictions are present. When this is the case, both colleges will implement matching frictions without necessarily coordinating on them.

This paper highlights that admission uncertainties and application limits can make both high-ranked and low-ranked colleges better off in their competition to attract desirable students. They might help low-ranked colleges attract strong applicants away from their high-ranked rivals, by offering better admission prospects. In turn, high-ranked colleges also benefit despite (and, due to) losing some desirable students, as this mitigates the competition for the best applicants. This model implies that a consensus towards centralized, stable mechanisms might be impeded not only because they leave the low-ranked college worse off; but also because the high-ranked college would suffer as a result of the staunch competition this would introduce.

3.7 Omitted Results

It will be useful to introduce some notation for omitted results and proofs. As we assumed $\alpha > \beta$ throughout our discussion, Σ_A^{hi} will denote the expected contribution to human capital of the students the high-ranked college gets in the stable matching, or in the equilibrium of the Simultaneous Applications game when the quality differential is high. Other variables introduced below follow similar notation.

$$\Sigma_A^{hi} = \Sigma_A^s := p \times \int_{1-\kappa_L}^{1} a.da + (1-p) \times \int_{1-\kappa_H}^{1} a.da$$

$$\Sigma_B^s := p \times \int_{1-2\kappa_L}^{1-\kappa_L} a.da + (1-p) \times \int_{1-2\kappa_H}^{1-\kappa_H} a.da$$

$$\Sigma_B^{hi} := p \times \int_{1-\kappa_L-\kappa_H}^{1-\kappa_H} a.da + (1-p) \times \int_{1-2\kappa_H}^{1-\kappa_H} a.da$$

$$\Sigma_A^{lo} := p \times \int_{1-\kappa_L}^{1} a.da + (1-p) \times \left(\int_{1-\kappa_L-\kappa_H}^{1-2\kappa_L} a.da + \int_{1-\kappa_L}^{1} a.da \right)$$

$$\Sigma_B^{lo} := p \times \int_{1-2\kappa_L}^{1-\kappa_L} a.da + (1-p) \times \left(\int_{1-2\kappa_L}^{1-\kappa_L} a.da + \int_{1-2\kappa_H}^{1-\kappa_L-\kappa_H} a.da \right)$$

Claim 3.1. The investment level α^{s*}, as defined in Section 3.3, is given by:

$$\alpha^{s*} := \frac{\Sigma_A^s + \sqrt{(\Sigma_A^s)^2 - (\Sigma_B^s)^2}}{2}$$

Proof. Note that $\Pi_A^s(\cdot)$ is continuously differentiable. For all $\alpha > \frac{\Sigma_A^s}{2}$, ${\Pi_A^s}'(\alpha) < 0$ and $\Pi_A^s(\alpha^{s*}) = \overline{\Pi}_B^s$. This proves the claim. □

The technical results Lemma 3.7, Lemma 3.8, Corollary 3.9, and Lemma 3.10 establish certain properties of B's constrained payoff functions. These are used to establish Proposition 3.5.

Lemma 3.7. The constrained conditional payoff functions $\widetilde{\Pi_A^{hi}}(\cdot)$, $\widetilde{\Pi_A^{lo}}(\cdot)$, $\widetilde{\Pi_B^{hi}}(\cdot)$, $\widetilde{\Pi_B^{lo}}(\cdot)$ are pairwise single-crossing.

Proof. I prove this statement for each pairwise combination. First take the combination of $\widetilde{\Pi}_B^{hi}$ with any other constrained conditional payoff function $\widetilde{\Pi}$. $\widetilde{\Pi}_B^{hi}$ is a constant function. Any other constrained conditional payoff function $\widetilde{\Pi}$ is weakly monotone, has $\lim_{\alpha\to\infty}\widetilde{\Pi} = -\infty$, and $\widetilde{\Pi}(0) > \widetilde{\Pi}_B^{hi}$; so, it crosses $\widetilde{\Pi}_B^{hi}$ exactly once. The crossing points of each respective curve is then identified as follows:

$$\widetilde{\alpha}_{B,hi}^{A,hi} := \frac{1-p}{2} \times (\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{hi})^2}) \tag{3.7.1}$$

$$\widetilde{\alpha}_{B,hi}^{A,lo} := \frac{(\Sigma_A^{lo} + \sqrt{(\Sigma_A^{lo})^2 - (\Sigma_B^{hi})^2})}{2} \tag{3.7.2}$$

$$\widetilde{\alpha}_{B,hi}^{B,lo} := \frac{(\Sigma_B^{lo} + \sqrt{(\Sigma_B^{lo})^2 - (\Sigma_B^{hi})^2})}{2(1-p)} \tag{3.7.3}$$

This leaves us with pairwise combinations of $\left\{\widetilde{\Pi}_A^{hi}, \widetilde{\Pi}_A^{lo}, \widetilde{\Pi}_B^{lo}\right\}$. Note that each of these curves are continuous.

1. $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$

   (i) Take the "parent curves" $\Pi_A^{hi}(\frac{\alpha}{1-p})$ and $\Pi_A^{lo}(\alpha)$, which, to the right of their peak, characterize these constrained conditional payoff functions. They cross at $\alpha = 0$ and $\alpha = \alpha_x := (1-p)\frac{\Sigma_A^{hi} - \Sigma_A^{lo}(1-p)}{p(2-p)}$. $\Pi_A^{hi}(\frac{\alpha}{1-p})$ peaks at $\hat{\alpha} := (1-p)\frac{\Sigma_A^{hi}}{2}$. Albeit tedious, it's straightforward to verify that for all $p \in (0,1)$, $\alpha_x > \hat{\alpha}$. So, these two curves cross after the point where the original curve Π_A^{hi} peaks.

   (ii) Suppose the crossing point α_x of the parent curves is to the left of $\Pi_A^{lo}(\alpha)$'s peak: $\alpha_x < \frac{\Sigma_A^{lo}}{2}$. As previously established, the curve $\widetilde{\Pi}_A^{hi}$ is strictly decreasing at this point. However, since $\alpha_x < \frac{\Sigma_A^{lo}}{2}$, the curve $\widetilde{\Pi}_A^{lo}(\alpha)$ is constant at α_x. Then, by the Intermediate Value Theorem, $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$ cross at a point to the left of α_x. They can cross at no other point: $\widetilde{\Pi}_A^{hi}$ decreases weakly faster than $\widetilde{\Pi}_A^{lo}$ everywhere. Thus, $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$ cross only at the point $\widetilde{\alpha}_{A,lo}^{A,hi}$ where $\Pi_A^{hi}(\frac{\alpha}{1-p})$ crosses $\widetilde{\Pi}_A^{lo}$'s peak:

   $$\widetilde{\alpha}_{A,lo}^{A,hi} := \frac{(1-p)}{2}(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_A^{lo})^2})$$

   (iii) Suppose instead that $\alpha_x \geq \frac{\Sigma_A^{lo}}{2}$. Then, I argue that this is the unique point $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$ cross.

Since the point α_x is to the right of the parent curves' peaks, the curves $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$ also cross here. If there is another crossing point, this must be in the region where $\widetilde{\Pi}_A^{lo}$ is constant. At this point, the first derivative of $\widetilde{\Pi}_A^{hi}$ is negative, and the first derivative of $\widetilde{\Pi}_A^{lo}$ is zero. However, it can be shown easily that the first derivative of $\widetilde{\Pi}_A^{hi}$ always decreases faster than that of $\widetilde{\Pi}_A^{lo}$, which means that a crossing at α_x cannot have happened.

This identifies the crossing point of these two curves, $\widetilde{\alpha}_{A,lo}^{A,hi}$, as:

$$\widetilde{\alpha}_{A,lo}^{A,hi} := \begin{cases} (1-p) \times \frac{\Sigma_A^{hi} - \Sigma_A^{lo}(1-p)}{p(2-p)} & (1-p) \times \frac{\Sigma_A^{hi} - \Sigma_A^{lo}(1-p)}{p(2-p)} \geq \frac{\Sigma_A^{lo}}{2} \\ \frac{(1-p)}{2}(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_A^{lo})^2}) & \text{otherwise} \end{cases} \tag{3.7.4}$$

2. $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_B^{lo}$

   (i) Again, consider the parent curves $\Pi_A^{hi}(\frac{\alpha}{1-p})$ and $\Pi_B^{lo}((1-p) \times \alpha)$. These curves cross at $\alpha = 0$ and $\alpha = \alpha_x := \dfrac{\frac{\Sigma_A^{hi}}{1-p} - (1-p)\Sigma_B^{lo}}{\frac{1}{(1-p)^2} - (1-p)^2}$. Again, tedious but straightforward algebra verifies that $\alpha_x > \frac{\Sigma_A^{hi}(1-p)}{2}$ for all $p > 0$. So, the parent curves cross to the right of where $\Pi_A^{hi}(\frac{\alpha}{1-p})$ peaks.

   (ii) Now, suppose $\alpha_x < \frac{\Sigma_B^{lo}}{2(1-p)}$; so that the parent curves cross before $\Pi_B^{lo}((1-p)\times\alpha)$ peaks. Again by the IVT, the crossing point of $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_B^{lo}$ must then lie to the left of α_x. Due to a similar argument we made for the curves $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$, this must be the only point they cross. Hence, $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_B^{lo}$ cross only at the point $\widetilde{\alpha}_{B,lo}^{A,hi}$ where the curve $\Pi_A^{hi}(\frac{\alpha}{1-p})$ crosses $\widetilde{\Pi}_B^{lo}$'s peak:

$$\widetilde{\alpha}_{B,lo}^{A,hi} = \frac{(1-p)}{2} \times \left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right)$$

   (iii) Suppose instead that $\alpha_x \geq \frac{\Sigma_B^{lo}}{2(1-p)}$. Again due to a similar argument we made for the curves $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_A^{lo}$, α_x is then the unique crossing point of the curves $\widetilde{\Pi}_A^{hi}$ and $\widetilde{\Pi}_B^{lo}$. This is where $\widetilde{\Pi}_A^{hi}$ crosses the constant portion of the curve $\widetilde{\Pi}_B^{lo}$:

$$\widetilde{\alpha}_{B,lo}^{A,hi} = \frac{\frac{\Sigma_A^{hi}}{1-p} - (1-p)\Sigma_B^{lo}}{\frac{1}{(1-p)^2} - (1-p)^2}$$

This identifies the crossing point of these two curves, $\widetilde{\alpha}^{A,hi}_{B,lo}$, as:

$$\widetilde{\alpha}^{A,hi}_{B,lo} = \begin{cases} \dfrac{\frac{\Sigma^{hi}_A}{1-p} - (1-p)\Sigma^{lo}_B}{\frac{1}{(1-p)^2} - (1-p)^2} & \frac{\frac{\Sigma^{hi}_A}{1-p} - (1-p)\Sigma^{lo}_B}{\frac{1}{(1-p)^2} - (1-p)^2} \geq \frac{\Sigma^{lo}_B}{2(1-p)} \\ \frac{(1-p)}{2} \times \left(\Sigma^{hi}_A + \sqrt{(\Sigma^{hi}_A)^2 - \left(\Sigma^{lo}_B\right)^2}\right) & \text{otherwise} \end{cases} \tag{3.7.5}$$

3. $\widetilde{\Pi}^{lo}_A$ and $\widetilde{\Pi}^{lo}_B$

   (i) Again, consider the parent curves $\Pi^{lo}_A(\alpha)$ and $\Pi^{lo}_B((1-p)\times\alpha)$. They cross at $\alpha = 0$ and $\alpha = \alpha_x := \frac{\Sigma^{lo}_A - (1-p)\Sigma^{lo}_B}{p(2-p)}$. Again, it is straightforwardly (but tediously) verifiable that $\alpha_x > \frac{\Sigma^{lo}_A}{2}$ for all $p > 0$.

   (ii) Suppose $\alpha_x < \frac{\Sigma^{lo}_B}{2(1-p)}$; so that the parent curves cross to the left of where $\Pi^{lo}_B((1-p)\times\alpha)$ peaks. Then, an argument mirroring the previous two pairs of curves establishes that the curves $\widetilde{\Pi}^{lo}_A$ and $\widetilde{\Pi}^{lo}_B$ must cross at the point $\widetilde{\alpha}^{A,lo}_{B,lo}$ where:

$$\widetilde{\alpha}^{A,lo}_{B,lo} := \frac{1}{2}\left(\Sigma^{lo}_A + \sqrt{\left(\Sigma^{lo}_A\right)^2 - \left(\Sigma^{lo}_B\right)^2}\right)$$

   (iii) Suppose instead $\alpha_x \geq \frac{\Sigma^{lo}_B}{2(1-p)}$. Again a similar argument to the two previous subcases then establish that α_x must then be the unique crossing point of the curves $\widetilde{\Pi}^{lo}_A$ and $\widetilde{\Pi}^{lo}_B$.

This identifies the crossing point $\widetilde{\alpha}^{A,lo}_{B,lo}$ of these two curves as:

$$\widetilde{\alpha}^{A,lo}_{B,lo} := \begin{cases} \frac{\Sigma^{lo}_A - (1-p)\Sigma^{lo}_B}{p(2-p)} & \frac{\Sigma^{lo}_A - (1-p)\Sigma^{lo}_B}{p(2-p)} \geq \frac{\Sigma^{lo}_B}{2(1-p)} \\ \frac{(\Sigma^{lo}_A + \sqrt{(\Sigma^{lo}_A)^2 - (\Sigma^{lo}_B)^2})}{2} & \text{otherwise} \end{cases} \tag{3.7.6}$$

□

Lemma 3.8. There exists some $\bar{p} \in (0,1)$ such that for all $p > \bar{p}$, the crossing points of B's constrained conditional payoff functions satisfy the two conditions:

$$\widetilde{\alpha}^{A,hi}_{A,lo} < \bar{\alpha}^{lo}_A \qquad\qquad \widetilde{\alpha}^{A,lo}_{B,lo} < \frac{\bar{\alpha}^{lo}_B}{1-p}$$

Furthermore, for all $p > \bar{p}$ and $\alpha \geq \widetilde{\alpha}^{B,lo}_{B,hi}$, $\Pi^{hi}_A(\alpha) < \bar{\Pi}^{lo}_B$.

Proof. Recall that $\bar{\alpha}_A^{lo} = \frac{\Sigma_A^{lo}}{2}$, and $\bar{\alpha}_B^{lo} = \frac{\Sigma_B^{lo}}{2}$ are the corresponding optimal investment levels.

I show that for large enough p, each individual claim holds:

(i) Consider $\widetilde{\alpha}_{A,lo}^{A,hi}$, identified in equation 3.7.4. Note that Σ_A^{hi} and Σ_A^{lo} are continuous and bounded functions of $p \in (0,1)$. Therefore, $\lim_{p\to 1}(1-p) \times \frac{\Sigma_A^{hi} - \Sigma_A^{lo}(1-p)}{p(2-p)} = 0$. As $\frac{\Sigma_A^{lo}}{2}$ is bounded below by $\Sigma_{A,\kappa_L}^{hi} > 0$, by continuity there is some $\bar{p_1} \in (0,1)$ such that for all $p > \bar{p_1}$, $(1-p) \times \frac{\Sigma_A^{hi} - \Sigma_A^{lo}(1-p)}{p(2-p)} < \frac{\Sigma_A^{lo}}{2}$ and $\widetilde{\alpha}_{A,lo}^{A,hi} = \frac{(1-p)}{2}\left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_A^{lo})^2}\right)$. Again, as Σ_A^{hi} and Σ_A^{lo} are continuous and bounded by positive real numbers for $p \in (0,1)$, there is some $\bar{p_2} \in (0,1)$ such that $p > \bar{p_2}$ implies $\widetilde{\alpha}_{A,lo}^{A,hi} < \frac{\Sigma_A^{lo}}{2}$. This proves the first claim; for $p > \max\{\bar{p_1}, \bar{p_2}\}$, $\widetilde{\alpha}_{A,lo}^{A,hi} < \frac{\Sigma_A^{lo}}{2} = \bar{\alpha}_A^{lo}$.

(ii) Now consider $\widetilde{\alpha}_{B,lo}^{A,lo}$, identified in equation 3.7.6. The inequality $\frac{\frac{\Sigma_A^{hi}}{1-p} - (1-p).\Sigma_B^{lo}}{\frac{1}{(1-p)^2} - (1-p)^2} \geq \frac{\Sigma_B^{lo}}{2(1-p)}$ can be written equivalently as $\frac{\Sigma_A^{hi}.(1-p) - (1-p)^3.\Sigma_B^{lo}}{1-(1-p)^4} \geq \frac{\Sigma_B^{lo}}{2(1-p)}$. Due to similar continuity and boundedness assumptions, as $p \to 1$, the LHS tends to 0, and the RHS to ∞. Similar to (i), there exists some $\bar{p_3} < 1$ such that $p > \bar{p_3}$ implies $\frac{\frac{\Sigma_A^{hi}}{1-p} - (1-p).\Sigma_B^{lo}}{\frac{1}{(1-p)^2} - (1-p)^2} < \frac{\Sigma_B^{lo}}{2(1-p)}$ then. Subsequently, for $p > \bar{p_3}$, $\widetilde{\alpha}_{B,lo}^{A,hi} = \frac{(1-p)}{2}.\left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right)$. But then, again similarly to (i), for some $\bar{p_4} < 1$, $p > \bar{p_4}$ implies that $\frac{(1-p)}{2}.\left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right) < \frac{\Sigma_B^{lo}}{2(1-p)} = \frac{\bar{\alpha}_B^{lo}}{1-p}$. This proves our second claim.

(iii) $\Pi_A^{hi}(\alpha) \geq \bar{\Pi}_B^{lo} \iff \alpha \in \left[\frac{1}{2}\left(\Sigma_A^{hi} - \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right), \frac{1}{2}\left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right)\right]$. Note, from equation 3.7.3, that since Σ_B^{lo} and Σ_B^{hi} are bounded for $p \in (0,1)$, $\lim_{p\to 1} \widetilde{\alpha}_{B,hi}^{B,lo} = \infty$, but $\lim_{p\to 1} \frac{1}{2}\left(\Sigma_A^{hi} + \sqrt{(\Sigma_A^{hi})^2 - (\Sigma_B^{lo})^2}\right)$ is bounded above. Since $\widetilde{\alpha}_{B,hi}^{B,lo}$ is continuous for $p < 1$, there exists some $\bar{p_5} < 1$ such that $p > \bar{p_5}$ implies $\Pi_A^{hi}(\alpha) < \bar{\Pi}_B^{lo}$ when $\alpha \geq \widetilde{\alpha}_{B,hi}^{B,lo}$.

For $\bar{p} := \max\{\bar{p_1}, \bar{p_2}, \bar{p_3}, \bar{p_4}, \bar{p_5}\}$, subcases (i), (ii), and (iii) hold when $p > \bar{p}$. This concludes the proof.

□

Corollary 3.9. Whenever Proposition 3.8 holds, the crossing points of B's constrained conditional payoff functions satisfy the following conditions:

$$\widetilde{\alpha}_{A,lo}^{A,hi} < \widetilde{\alpha}_{B,lo}^{A,hi} \qquad \widetilde{\alpha}_{A,lo}^{A,hi} < \widetilde{\alpha}_{B,hi}^{A,hi} \qquad \widetilde{\alpha}_{B,lo}^{A,lo} < \widetilde{\alpha}_{B,hi}^{A,lo}$$

Proof. $\widetilde{\alpha}^{A,hi}_{A,lo} < \bar{\alpha}^{lo}_{A}$ implies that $\widetilde{\Pi}^{lo}_{A}(\widetilde{\alpha}^{A,hi}_{A,lo}) = \bar{\Pi}^{lo}_{A} = \widetilde{\Pi}^{hi}_{A}(\widetilde{\alpha}^{A,hi}_{A,lo})$. Since for all $\alpha \geq 0$, $\widetilde{\Pi}^{lo}_{B}(\alpha) < \bar{\Pi}^{lo}_{A}$, $\widetilde{\Pi}^{lo}_{B}(\widetilde{\alpha}^{A,hi}_{B,lo}) = \widetilde{\Pi}^{hi}_{A}(\widetilde{\alpha}^{A,hi}_{B,lo}) < \bar{\Pi}^{lo}_{A}$. As $\widetilde{\Pi}^{hi}_{A}(\alpha)$ is monotone non-increasing, this implies that $\widetilde{\alpha}^{A,hi}_{A,lo} < \widetilde{\alpha}^{A,hi}_{B,lo}$. Similarly, $\widetilde{\Pi}^{hi}_{A}(\widetilde{\alpha}^{A,hi}_{B,hi}) < \bar{\Pi}^{lo}_{A}$. This implies that $\widetilde{\alpha}^{A,hi}_{A,lo} < \widetilde{\alpha}^{A,hi}_{B,hi}$.

$\widetilde{\alpha}^{A,lo}_{B,lo} < \frac{\bar{\alpha}^{lo}_{B}}{1-p}$implies that $\widetilde{\Pi}^{lo}_{B}(\widetilde{\alpha}^{A,lo}_{B,lo}) = \bar{\Pi}^{lo}_{B} = \widetilde{\Pi}^{lo}_{A}(\widetilde{\alpha}^{A,lo}_{B,lo})$. As $\widetilde{\Pi}^{lo}_{A}(\widetilde{\alpha}^{A,lo}_{B,hi}) = \bar{\Pi}^{hi}_{B}$, it similarly follows that $\widetilde{\alpha}^{A,lo}_{B,lo} < \widetilde{\alpha}^{A,lo}_{B,hi}$. □

Lemma 3.10. Suppose that at some fixed investment level α for A, one constrained conditional payoff curve is strictly above the rest. Then, B's best response to this investment level by A involves attracting the corresponding set of applicants.

Proof. For given investment levels α and β for A and B, B attracts the following set of applicants:

$$
\begin{aligned}
&(\mu^{hi})^{-1}(A,\kappa) && \text{for } \beta > \frac{\alpha}{1-p} \\
&(\mu^{lo})^{-1}(A,\kappa) && \text{for } \frac{\alpha}{1-p} > \beta > \alpha \\
&(\mu^{lo})^{-1}(B,\kappa) && \text{for } \alpha > \beta > \alpha \times (1-p) \\
&(\mu^{hi})^{-1}(B,\kappa) && \text{for } \alpha \times (1-p) > \beta
\end{aligned}
$$

Now, fix a set of applicants. The constrained conditional payoff function for that set of applicants, evaluated at A's investment level α, is the supremum of the expected payoff B can achieve with those applicants given β weakly exceeds the level of investment required to attract those applicants. To see this, fix the set of applicants $(\mu^{hi})^{-1}(A,\kappa)$, whose constrained conditional payoff function is $\widetilde{\Pi^{hi}_{A}}(\cdot)$. Letting α be A's chosen investment level:

- if $\alpha^{\bar{hi}} > \frac{\alpha}{1-p}$, then the investment level $\beta = \alpha^{\bar{hi}}$ is weakly higher than needed to secure this applicant group and attains the payoff $\widetilde{\Pi^{hi}_{A}}(\alpha) = \bar{\Pi}^{hi}_{A}$.
- if $\alpha^{\bar{hi}} \leq \frac{\alpha}{1-p}$, college B can obtain an expected payoff arbitrarily close to $\Pi^{hi}_{A}\left(\frac{\alpha}{1-p}\right)$ by setting $\beta = \frac{\alpha}{1-p} + \varepsilon$ for an arbitrarily small $\varepsilon > 0$.

The reasoning is analogous for the remaining sets of applicants.

Now, take a fixed level of investment α. Suppose that, at this point α, one of the four constrained conditional payoff curves exceeds all others. Denote the corresponding set of applicants as μ^{-1}. I now prove that college B's best response involves attracting

this set of applicants. The argument above clearly establishes that given A's investment α, B cannot attract any other set of applicants at higher payoff. It also establishes that, at any investment level β for which the expected payoff of B given the set of applicants μ^{-1} is arbitrarily close to the value of the constrained conditional payoff function for these applicants, B attracts either the set of applicants μ^{-1}, or a weakly better set of applicants $(\mu')^{-1}$. However, B cannot attract a strictly better set of applicants $(\mu')^{-1}$ at such investment levels β: if it could, the conditional constrained payoff function for the applicants $(\mu')^{-1}$ would need to lie strictly above that of $(\mu)^{-1}$, since at the same investment level, the applicants $(\mu')^{-1}$ must yield greater expected payoff than the applicants $(\mu)^{-1}$. □

3.8 Omitted Proofs

Proof, Lemma 3.1: Following Gale and Shapley (1962) and Abdulkadiroğlu, Che, and Yasuda (2015), the Deferred Acceptance algorithm produces a stable matching. Therefore, if $\mu^s(a,\kappa)$ is the unique stable matching of the college admissions game, this must be the realised matching.

Now, let us prove that $\mu^s(a,\kappa)$ is the unique stable matching in the college admissions game. Clearly, $\mu^s(a,\kappa)$ is individually rational. Given realised capacities κ, A strictly prefers any student it is matched to over any other student. So, no pairwise block can involve A. Any pairwise block involving B must involve a student with type $a \geq 1-\kappa$. However these students are matched to A, which they strictly prefer over B. Hence, $\mu^s(a,\kappa)$ is stable.

Now consider any other stable matching $\mu(a,\kappa) \neq \mu^s(a,\kappa)$. Since all students are acceptable and preferences are responsive, no matching where some college has unsaturated capacity can be stable, so restrict attention to matchings where colleges' capacities are saturated. $\mu(a,\kappa)$'s stability implies $\mu(A,\kappa) = \mu^s(A,\kappa)$; otherwise A would pairwise block μ with any student of type a such that $\mu(a,\kappa) \neq A$ but $\mu^s(a,\kappa) = A$.

But then, $\mu(B,\kappa) = \mu^s(B,\kappa)$ also; since students matched to A are unattainable for B, and any student with type a such that $\mu^s(a,\kappa) = B$ is strictly preferable over any with type $\mu^s(a,\kappa) = \emptyset$. $\mu = \mu^s$ therefore; yielding a contradiction. □

Proof, Proposition 3.2: First, note that setting $\beta = \alpha$ is a strictly dominated strategy for college B. When $\alpha = \beta$, students have no strict preference for either college. Therefore, they will choose which college to prioritise in their applications at random. Thus, B will fail to enrol the full mass of students for whom $a \in [\kappa, 1]$. Instead, by deviating to $\beta = \alpha + \varepsilon$ for some arbitrarily small $\varepsilon > 0$, B can guarantee enrolling the full mass of students for whom $a \in [\kappa, 1]$. It is easily seen that by choosing $\varepsilon > 0$ small enough, college B can guarantee a strict payoff gain from this deviation.

For any $\alpha < \alpha^{s*}$, B can set $\beta = \alpha + \varepsilon$ for an arbitrarily small $\varepsilon > 0$ to obtain a payoff above $\bar{\Pi}^s_B$. This leaves A a payoff of $\Pi^s_B(\alpha)$. Therefore, only the strategy $\alpha = \bar{\beta}^s$ can be optimal among the strategies $\alpha < \alpha^{s*}$. This strategy yields an expected payoff of $\bar{\Pi}^s_B$ for A.

For $\alpha \geq \alpha^{s*}$, $\beta > \alpha$ yields B an expected payoff weakly below $\bar{\Pi}^s_B$. Therefore, B's best response is to set $\beta = \bar{\beta}^s$. A then obtains the expected payoff $\Pi^s_A(\alpha)$. For $\alpha \geq \alpha^{s*}$, this is maximised at $\alpha = \alpha^{s*}$, and $\Pi^s_A(\alpha^{s*}) = \bar{\Pi}^s_B$.

□

Proof, Proposition 3.3: I first prove that μ^{hi} is the unique matching that can emerge as an equilibrium outcome of this game. To do so, I use the Iterated Elimination of Strictly Dominated Strategies (IESDS).

Consider students with types $a \geq 1 - \kappa_L$. No college can ever form a cohort of measure κ_L with students of types $a' > a$; so, $\rho_A(a, \sigma) = \rho_B(a, \sigma) = 1$ for any strategy profile $\sigma(\cdot) : [0, 1] \to [0, 1]$ students may employ. Thus, for all $a \geq 1 - \kappa_L$, the equilibrium strategies σ^{hi*} must be such that for all $a \geq 1 - \kappa_L$, $\sigma^{hi*}(a) = 1$.

Now, consider students with types $a \in [1 - \kappa_H, 1 - \kappa_L)$. No college can form a cohort of measure κ_H with students of types $a' > a$; so, $\rho_A(a, \sigma) \geq 1 - p$ for any strategy profile $\sigma(\cdot) : [0, 1] \to [0, 1]$ students may employ. This implies that $\frac{\alpha}{\beta} > \frac{\rho_B(a,\sigma)}{\rho_A(a,\sigma)}$ for any such strategy profile; so, the equilibrium strategies must be such that for all $a \in [1-\kappa_H, 1-\kappa_L)$, $\sigma^{hi*}(a) = 1$. The remaining students thus have a probability 0 of enrolling at college A, and without loss, we may assume that $\sigma^{hi*}(a) = 0$ for such students. This generates the matching μ^{hi} as claimed.

The matching $\mu^{hi}(a, \kappa_H)$ coincides with $\mu^s(a, \kappa_H)$, and is therefore stable. The match-

ing $\mu^{hi}(a, \kappa_L)$ differs from $\mu^s(a, \kappa_L)$; so, by Lemma 3.1, it exhibits justified envy.

□

Proof, Proposition 3.4: I first prove that μ^{lo} is the unique matching that can emerge as an equilibrium outcome of this game. To do so, again by IESDS.

Denote students' equilibrium application strategies in this game as σ^{lo*}. As established in Proposition 3.3's proof, $\sigma^{lo*}(a) = 1$ for all $a \geq 1 - \kappa_L$. By IESDS, we then have $\rho_A(a, \sigma^{lo*}) \leq (1 - p)$ for all $a \in [1 - 2\kappa_L, 1 - \kappa_L)$. Moreover, B cannot form a cohort of size κ_L from students with $a' > a$ for any $a \in [1 - 2\kappa_L, 1 - \kappa_L)$. Therefore, $\rho_B(a, \sigma^{lo*}) = 1$ for all $a \in [1 - 2\kappa_L, 1 - \kappa_L)$. Since this implies that $\frac{\alpha}{\beta} < \frac{\rho_B(a,\sigma^{lo*})}{\rho_A(a,\sigma^{lo*})}$, we have $\sigma^{lo*}(a) = 0$ for all $a \in [1 - 2\kappa_L, 1 - \kappa_L)$.

Now, consider students with $a \in [1 - \kappa_L - \kappa_H, 1 - 2\kappa_L)$. Given the equilibrium strategies of better students, no college admits these students when $\kappa = \kappa_L$. However, both colleges will admit them when $\kappa = \kappa_H$. So, these students will apply to college A in equilibrium; i.e., $\sigma^{lo*}(a) = 1$ for all such a.

Finally, consider students with $a \in [1 - 2\kappa_H, 1 - \kappa_L - \kappa_H)$. Given the equilibrium strategies of better students, these students face no chance of enrolling at A. Therefore, they all apply to college B in equilibrium; $\sigma^{lo*}(a) = 0$ for all such a.

These equilibrium strategies give rise to the matching μ^{lo} in equilibrium. $\mu^{lo}(a, \kappa_L) = \mu^s(a, \kappa_L)$; so, this matching is stable when $\kappa = \kappa_L$. However, $\mu^{lo}(a, \kappa_H) \neq \mu^s(a, \kappa_H)$; so, this matching exhibits justified envy when $\kappa = \kappa_H$.

□

Proof, Proposition 3.5: Suppose that for $p > \bar{p}$, Proposition 3.8 and Corollary 3.9 hold. Then, Lemma 3.10 asserts that B's set of preferred applicants, given A's investment level α, are:

$$
\begin{aligned}
&(\mu^{hi})^{-1}(B, \kappa) && \text{for } \alpha \geq \tilde{\alpha}^{B,lo}_{B,hi} \\
&(\mu^{lo})^{-1}(B, \kappa) && \text{for } \tilde{\alpha}^{B,lo}_{B,hi} > \alpha \geq \tilde{\alpha}^{A,lo}_{B,lo} \\
&(\mu^{lo})^{-1}(A, \kappa) && \text{for } \tilde{\alpha}^{A,lo}_{B,lo} > \alpha \geq \tilde{\alpha}^{A,hi}_{A,lo} \\
&(\mu^{hi})^{-1}(A, \kappa) && \text{for } \tilde{\alpha}^{A,hi}_{A,lo} > \alpha
\end{aligned}
$$

The set of applicants A attracts at the same level of investment is then given by:

$$\begin{aligned}
&(\mu^{hi})^{-1}(A,\kappa) && \text{for } \alpha \geq \tilde{\alpha}^{B,lo}_{B,hi} \\
&(\mu^{lo})^{-1}(A,\kappa) && \text{for } \tilde{\alpha}^{B,lo}_{B,hi} > \alpha \geq \tilde{\alpha}^{A,lo}_{B,lo} \\
&(\mu^{lo})^{-1}(B,\kappa) && \text{for } \tilde{\alpha}^{A,lo}_{B,lo} > \alpha \geq \tilde{\alpha}^{A,hi}_{A,lo} \\
&(\mu^{hi})^{-1}(B,\kappa) && \text{for } \tilde{\alpha}^{A,hi}_{A,lo} > \alpha
\end{aligned}$$

By Proposition 3.8, $\Pi^{lo}_A(\widetilde{\alpha}^{A,lo}_{B,lo}) = \bar{\Pi}^{lo}_B$. When $\alpha = \widetilde{\alpha}^{A,lo}_{B,lo}$, B prefers settling for the applicants $(\mu^{lo})^{-1}(B,\kappa_H)$ rather than competing for $(\mu^{lo})^{-1}(A,\kappa_H)$. Therefore, $\Pi_A(\widetilde{\alpha}^{A,lo}_{B,lo}) = \bar{\Pi}^{lo}_B$.

Setting $\alpha \in (\widetilde{\alpha}^{A,lo}_{B,lo}, \widetilde{\alpha}^{B,lo}_{B,hi})$ cannot be optimal, as $\Pi_A(\alpha) = \Pi^{lo}_A(\alpha) < \bar{\Pi}^{lo}_B$ for such α. The last part of Proposition 3.8 asserts that setting $\alpha \geq \widetilde{\alpha}^{B,lo}_{B,hi}$ cannot be optimal either, for $\alpha \geq \widetilde{\alpha}^{B,lo}_{B,hi}$ implies $\Pi_A = \Pi^{hi}_A(\alpha) < \bar{\Pi}^{lo}_B$. For $\alpha < \widetilde{\alpha}^{A,lo}_{B,lo}$, A can at most get an equilibrium payoff of $\bar{\Pi}^{lo}_B$ since it will get students weakly inferior to $(\mu^{lo})^{-1}(B,\kappa_H)$. As colleges break ties in favour of better students, A's equilibrium investment strategy α^* is then $\alpha^* = \widetilde{\alpha}^{A,lo}_{B,lo}$.

By Proposition 3.8, $\widetilde{\alpha}^{A,lo}_{B,lo} < \frac{\bar{\alpha}^{lo}_B}{1-p}$. Then, college B's best response β^* to the investment level α^* set by college A is $\frac{\Sigma^{lo}_B}{2}$; this is the optimal level of investment $\bar{\alpha}^{lo}_B$ for the set of applicants $(\mu^{lo})^{-1}(B,\kappa_H)$.

As $\widetilde{\alpha}^{A,lo}_{B,lo} > \frac{\Sigma^{lo}_B}{2} > (1-p) \times \widetilde{\alpha}^{A,lo}_{B,lo}$, the strategy σ^{lo*} defined in the proof of Proposition 3.4 describes students' application behaviour. Then, colleges A and B get the applicants $(\mu^{lo})^{-1}(A,\kappa_H)$ and $(\mu^{lo})(B,\kappa_H)$ in equilibrium, respectively. They thus expect the payoff $\bar{\Pi}^{lo}_B$ in equilibrium.

□

Proof, Proposition 3.6. Let p be large enough so that Proposition 3.5 holds. Then, in Simultaneous Applications, both colleges equilibrium payoffs are $\bar{\Pi}^{lo}_B$. This exceeds their equilibrium payoff in the Sequential Applications game, where B secures a worse intake of students, in expectation. Therefore, in the equilibrium of the college admissions game, the college that sets its date second will always mirror the first college's date.

□

References

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda (Feb. 2015). "Expanding "Choice" in School Choice". In: *American Economic Journal: Microeconomics* 7.1, pp. 1–42. DOI: 10.1257/mic.20120027.

Avery, Christopher, Soohyung Lee, and Alvin E Roth (2014). *College Admissions as Non-Price Competition [electronic resource]: The Case of South Korea*. eng. Working paper series (National Bureau of Economic Research : Online) ; working paper no. 20774. Cambridge, Mass.: National Bureau of Economic Research.

Avery, Christopher and Jonathan Levin (Dec. 2010). "Early Admissions at Selective Colleges". In: *American Economic Review* 100.5, pp. 2125–56. DOI: 10.1257/aer.100.5.2125.

Barkhorn, Eleanor (Jan. 2014). "College Rankings Really Do Influence Which Schools Students Apply To". In: *The Atlantic*. Accessed: 2025-07-23.

Burdett, Ken and Melvyn G. Coles (2001). "Transplants and Implants: The Economics of Self-Improvement". In: *International Economic Review* 42.3, pp. 597–616. ISSN: 00206598, 14682354. DOI: 10.1111/1468-2354.00126.

Chade, Hector, Gregory Lewis, and Lones Smith (2014). "Student Portfolios and the College Admissions Problem". eng. In: *The Review of economic studies* 81.3 (288), pp. 971–1002. ISSN: 0034-6527. DOI: 10.1093/restud/rdu003.

Chen, Wei-Cheng and Yi-Cheng Kao (2014). "Simultaneous screening and college admissions". In: *Economics Letters* 122.2, pp. 296–298. DOI: 10.1016/j.econlet.2013.12.

— (2023). "Limiting Applications in College Admissions and Evidence from Conflicting Exam Dates". eng. In: *Journal of Human Capital* 17.3, pp. 434–461. ISSN: 1932-8575. DOI: 10.1086/724567.

Eom, Taekyu (Nov. 2024). "Cap-and-Apply: Unintended Consequences of Restricting College Applications". Paper presented at the ASSA 2025 Annual Meeting, San Francisco, CA. Manuscript dated November 11, 2024. URL: https://www.aeaweb.org/conference/2025/program/paper/R9D5dsFs (visited on 07/23/2025).

Felli, Leonardo and Kevin Roberts (2016). "Does Competition Solve the Hold-up Problem?" In: *Economica* 83.329, pp. 172–200. ISSN: 00130427, 14680335. DOI: https://doi.org/10.1111/ecca.12170.

Gale, D. and L. S. Shapley (1962). "College Admissions and the Stability of Marriage". In: *The American Mathematical Monthly* 69.1, pp. 9–15. ISSN: 00029890, 19300972. DOI: 10.2307/2312726.

Tierney, John (Sept. 2013). "Your Annual Reminder to Ignore the U.S. News & World Report College Rankings". In: *The Atlantic*. Accessed: 2025-07-23.

Trachtenberg, Stephen Joel (Sept. 2011). "How Dangerous Are College Rankings and the Rat Race for Prestige?" In: *The Atlantic*. Accessed: 2025-07-23.

University of Oxford (2025). *UCAS application*. https://www.ox.ac.uk/admissions/undergraduate/applying-to-oxford/guide/ucas-application. Accessed: 2025-07-23.

Conclusion

This PhD dissertation presents three stand-alone papers in the field of information economics.

Chapter 1 explores the link between efficiency and information in decentralised markets. It uncovers a trade-off: buyers need information to discern whether there are gains from trading with the seller, but are exposed to more adverse selection in a market where more information is available. How this trade-off is resolved depends not on the amount, but on the *kind* of information available in the market.

This chapter contributes to a large literature on information aggregation in markets. The bulk of this literature asks whether market outcomes fully reflect market participants' collective information; i.e., whether markets aggregate information efficiently. Instead, Chapter 1 asks whether market outcomes improve when market participants are collectively more informed; i.e., whether a decentralised market aggregates information *better* when more information is available.

This chapter highlights that the microstructure of a marketplace not only determines its ability to efficiently aggregate information, but also whether and when it becomes more efficient as more information becomes available. Its findings relate directly to policy; for instance, its main results align with recent ECB (2024) guidance on the use of different kinds of information in credit markets. This suggests that exploring the questions this chapter puts forward in other market contexts is a promising area of future research. It is also increasingly urgent as technological developments continue to dramatically ease access to information across financial markets.

Chapter 2 studies the optimal voting mechanism for a committee that must decide whether to enact or block a policy of unknown benefit. It finds that the dictatorship of the

most-demanding member is a dominant voting mechanism: any other voting mechanism is (i) less likely to enact a good policy, (ii) more likely to enact a bad policy, and (iii) burdens each member with a greater cost of acquiring information. However, it also points out that this sharp result may fail when certain assumptions in the model are relaxed; particularly:

- if information provision is no longer costless or unrestricted for the lobbyist, or
- if committee members can coordinate to equilibria which exhibit an *encouragement effect* in the spirit of Bolton and Harris (1999).

Chapter 2 does not explore whether relaxing the model as such preserves some desirable properties of the dictatorship of the most-demanding member, nor how the optimal voting mechanism changes as a result. These questions are of clear theoretical and policy interest and remain promising areas of future research.

Chapter 3 studies how application caps influence college admissions outcomes when students face uncertainty about their admission chances. It finds that such caps generate two effects. The *allocative effect* benefits the low-ranked college: it helps this college attract high-ability students who would have been admitted to the high-ranked college, but were unwilling to risk applying there. The *competitive effect* benefits the high-ranked college: since the low-ranked college's applicant pool improves, its incentive to compete for the very top students diminishes. These two effects incentivise *both* colleges to impose caps and costs on students' applications.

Chen and Kao (2023) find empirical evidence for the allocative effect: they document that in Taiwan, some departments of a second-rank college strategically align their admission dates with the top-ranked college to prevent students from applying to both, and that this improves the applicant pool of the second-rank college. Their study does not, however, analyse the competitive effect I identify. An empirical verification of this second effect remains a promising and policy-relevant area for future research.