Essays in Microeconomic Theory

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Justus Matthias Preusser

aus Essen

2024

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Stephan Lauermann
Zweitreferent:	Prof. Dr. Benny Moldovanu
Tag der mündlichen Prüfung:	26. Januar 2024

Acknowledgements

This thesis was made possible by the support of many people. I am deeply indebted to my advisors Stephan Lauermann, Benny Moldovanu and Francesc Dilmé for their early encouragement and continuing to believe in me. You supported me when I wanted to pursue my own ideas, but you nevertheless offered your advice and held me to the highest standards. For that I am very grateful.

I am immensely lucky to have found four fantastic co-authors Axel, Alex, Deniz, and Andre. Working with you has been a joy and this dissertation would not exist without you. I am especially grateful to my friend Axel. Our lively discussions and your sharp and honest comments were an integral part of these years.

For their friendship, I would like to express my heartfelt appreciation to all the wonderful people I have met on this journey, especially Valentin, Moritz, Paul, Thomas, Fabian, Sonja, Thorin, Raffael, Alessandro, Anusha, Sophie, Simon, and Kailin.

I am grateful to the great theory group at Bonn and all faculty members, students, postdocs, and staff members that helped me during the PhD and the job market. I have benefited from funding and the infrastructure of the BGSE, briq, CRC TR 224, and the German Academic Scholarship Foundation. I would also like to thank the University of California, Berkeley and Yale University for their hospitality.

I dedicate this thesis to my family, for their unconditional love and support.

Justus Matthias Preusser
Bonn, March 2023

Contents

List of Figures

Introduction

This thesis comprises three essays in microeconomic theory.

Chapter 1 is based on Niemeyer and Preusser (2022). We investigate the problem of allocating a scarce resource among a number of agents. The efficient allocation depends on the information that agents have about their peers, but monetary transfers are not used to elicit this information. This problem arises, for example, when distributing social grants among the members of a community (who can vouch for each other's needs), allocating funding to researchers (who can evaluate one other's work), or selecting the leader of the group (where each member has an opinion about the leadership qualities of their peers). We consider dominant-strategy incentive-compatible mechanisms for this problem. Our contribution is twofold. First, we establish fundamental properties of the set of mechanisms: deterministic mechanisms do not suffice for implementation, and anonymous mechanisms cannot meaningfully elicit information from the agents. Second, we propose and make a case for a simple class of mechanisms called *jury mechanisms*. These mechanisms solve the problem with three agents, are approximately optimal with many "exchangeable" agents, and are the only deterministic mechanisms satisfying a relaxed notion of anonymity.

Chapter 2 is based on Kattwinkel et al. (2022). We study a mechanism design problem with two agents having opposing interests. The leading example is an allocation problem. The management of a firm has to decide how to split a budget between two departments of the firm. The departments have private information about the marginal revenue they can generate, and management wishes to allocate to the department with the highest marginal revenue. As in Chapter 1, the departments may have valuable information about one another. This information is modelled as a type that correlates with the marginal revenues as well with the type of the other department. Our results shed light on how this information is useful to the manager. First, we fully characterize the set of implementable mechanisms. Further, in a sense that we make precise, we show that the manager is reliant on each department's type's being informative about the revenue of the *other* department. However, correlation between the types themselves shrinks the set of implementable mechanisms, revealing how the absence of transfers overturns canonical results due to Crémer and McLean (1985, 1988).

Chapter 3 is based on Preusser and Speit (2022). We study a dynamic market for a single good where market participants gradually learn about the good's value. For example, a prospective buyer of a house may update on the fact that other interested buyers inspected the house but chose not to buy. The precise inference depends on a number of details, such the house's time-on-the-market and the prices that earlier buyers were offered by the seller of the house. Consequently, market outcomes, too, depend on the transparency of the market. This paper contributes to a literature investigating how this form of transparency interacts with other features of the market (see, for example, Hörner and Vieille, 2009; Kim, 2017). We find that, when either everything or nothing is made public about the time-on-the-market and past prices, then the seller can appropriate all gains from trade. However, when time-on-the-market but not past prices are made public, buyers always enjoy gains from trade.

References

Crémer, Jacques, and Richard P. McLean. 1985. "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent." *Econometrica,* 345–61. [1]

Crémer, Jacques, and Richard P. McLean. 1988. "Full extraction of the surplus in Bayesian and dominant strategy auctions." *Econometrica,* 1247–57. [1]

Hörner, Johannes, and Nicolas Vieille. 2009. "Public vs. private offers in the market for lemons." *Econometrica* 77 (1): 29–69. [2]

Kattwinkel, Deniz, Axel Niemeyer, Justus Preusser, and Alexander Winter. 2022. "Mechanisms without transfers for fully biased agents." https://arxiv.org/abs/2205.10910. [1]

Kim, Kyungmin. 2017. "Information about sellers' past behavior in the market for lemons." *Journal of Economic Theory* 169: 365–99. [2]

Niemeyer, Axel, and Justus Preusser. 2022. "Simple Allocation with Correlated Types." https://jpreusser.github.io/. [1]

Preusser, Justus, and Andre Speit. 2022. "Transparency in Sequential Common-Value Trade." https://doi.org/https://jpreusser.github.io/. [2]

Chapter 1

Simple Allocation with Correlated Types

1.1 Introduction

We consider environments where an object is allocated among a number of agents. The efficient allocation depends on how the agents evaluate their peers, but monetary transfers are not used to elicit this information. A number of environments fit this description:

(1) A group has to elect one of its members to a prestigious post. The group as whole benefits from selecting a qualified candidate, and each agent knows the qualities of their friends in the group. Monetary transfers would naturally be excluded in such an election.

(2) A community of households has to distribute a good among its members. Each member can vouch for the needs and valuations of their friends or neighbors. If some members are financially constrained, it may be infeasible or undesirable to have members compete for the good via bids.

(3) A funding agency splits a budget across researchers. Each researcher can evaluate others in their field. If all parties are risk neutral, the allocated share of the budget can be interpreted as the probability of being allocated the object. Additional monetary transfers would be self-defeating.

In these environments, asking the agents straightforwardly who "should" get the object does not guarantee satisfactory outcomes. In particular, if agents are primarily concerned with their own winning chances, they may exaggerate their individual qualities instead of impartially disclosing their peer information.

To better understand good allocation rules, we take a mechanism design approach and consider the following model. Each agent wants to win the object and is indifferent to which of the others wins. Allocating to an agent generates a social value. The agents have private information about these values—their *types*. We model peer information by allowing for an arbitrary joint distribution of types and values. Hence an agent's type may be informative about the types and values of all others.

We study mechanisms for maximizing the expected value of the allocation. In a mechanism, each agent is asked to report their type. We focus on mechanisms where truthfully reporting one's type is a dominant strategy; that is, we focus on dominant-strategy incentive-compatible (DIC) mechanisms. For the assumed preferences of the agents, DIC requires that one's report never influences one's own winning probability.

Let us highlight some of the differences to existing models (a detailed review follows later). Alon et al. (2011) and Holzman and Moulin (2013) consider DIC mechanisms (there called strategyproof or impartial) where the agents nominate one another to win the object. These nominations do not arise from some ground truth. By contrast, we fix a general joint distribution of types and values. This lets us study mechanisms where, say, two agents can share their private information and form a consensus about which of the others to nominate. Other work considers settings where non-monetary instruments for screening the agents are available, but where the agents have no peer information (for example, Ben-Porath, Dekel, and Lipman, 2014, 2019).

We contribute two results demonstrating the difficulty of designing "simple" mechanisms for this problem: deterministic DIC mechanisms are not without loss, and anonymous DIC mechanism cannot meaningfully elicit information. We further contribute three positive results on so-called *jury mechanisms*. These mechanisms, described in detail below, solve the problem with three agents, are approximately optimal in symmetric environments with many agents, and are the only deterministic DIC mechanisms satisfying a relaxed notion of anonymity. Let us elaborate.

For each agent, there is a trade-off between allocating to the agent and using the agent's peer information. This trade-off arises since, on the one hand, DIC demands that a change in an agent's type does not affect that agent's own winning probability, but, on the other hand, the change in the type reveals information about the values from allocating to the others.

Optimally resolving this trade-off may require the use of stochastic mechanisms that cannot be implemented by randomizing over deterministic ones. That is, the set of DIC mechanisms may admit stochastic extreme points, and these can be uniquely optimal. Stochastic extreme points exist if and only if there are at least four agents and the type spaces are not "too small." The typical view in the literature is that one should use mechanisms that can be implemented by randomizing over deterministic ones (for example, Pycia and Ünver, 2015; Chen et al., 2019). We find that doing so is not generally without loss in the present problem.

Our next result is that all anonymous DIC mechanisms must ignore the reports of the agents. Here, anonymity means that all agents can make the same reports and that an agent's winning probability does not change when one permutes the reports of the others. We view anonymity as the familiar axiom from social choice theory that no agent play a special role in determining the chosen social alternative; that is, in determining who wins the object. As such, anonymity helps reduce the complexity of the mechanism, protects agents' privacy when evaluating their peers, and ensures that agents have the same rights as voters. Our negative result also sheds new light on a characterization due to Holzman and Moulin (2013) and Mackenzie (2015) of a slightly different notion of anonymity.

Our positive results concern the following class of mechanisms. In a *jury mechanism*, each agent is either a *juror* or a *candidate*. The allocation only depends on the reports of the jurors, and the object is always allocated to a candidate. Given that jurors cannot win, all jury mechanisms are DIC.

If there are three agents, then all DIC mechanisms are randomizations over deterministic jury mechanisms. In particular, a deterministic jury mechanism is optimal. This generalizes a known result for deterministic DIC mechanisms due to Holzman and Moulin (2013). Our key insight is that in the three-agent case all DIC mechanisms are actually randomizations over deterministic ones.

Next, we identify a condition on the environment under which deterministic jury mechanisms are approximately optimal with many agents. By "approximately optimal" we mean that the difference in expected values between an optimal deterministic jury mechanism and an optimal DIC mechanism vanishes as the number of agents diverges. The condition on the environment is that agents are exchangeable in terms of supplying information about the vector of values. Intuitively, when agents are exchangeable, increasing their number relaxes the aforementioned trade-off. In particular, there is essentially no loss from ignoring the reports of those agents who are sometimes allocated the object—this is the defining property of a jury mechanism.

For the last result, we consider a relaxed notion of anonymity—*partial anonymity*. Whereas the earlier notion of anonymity demands that an agent's winning probability be invariant with respect to *all* permutations of the others, partial anonymity only considers permutations of those agents that in the given mechanism *actually* influence the agent's winning probability. We show that all deterministic partially anonymous DIC mechanism are jury mechanisms.

The paper is organized as follows. We next discuss related work (Section 1.2) and present the model (Section 1.3). In Section 1.4, we introduce jury mechanisms and present the results for the three- and many-agent cases. In Section 1.5, we characterize when stochastic extreme points exist. In Section 1.6, we study anonymous mechanisms, presenting the two notions and the associated characterizations side-by-side. We conclude by discussing open questions (Section 1.7). All omitted proofs are in Appendix 1.A. Supplementary material is collected in Appendix 1.B and Appendix 1.C.

1.2 Related literature

Holzman and Moulin (2013) study axioms for peer nomination rules. In such a rule, agents nominate one another to receive a prize. Their central axiom—*impartiality*—is equivalent to DIC when each agent cares only about their own winning probability. As Holzman and Moulin note, many of their axioms have no obvious counterparts in a model with abstract types. Most relevant for us is their notion and characterization of anonymity, as well subsequent results due to Mackenzie (2015, 2020). We discuss the differences to our characterization in detail in Section 1.6.4.[1]

Alon et al. (2011) initiated a literature on optimal DIC mechanisms (there called *strategyproof* mechanisms) in a model where each agent nominates a subset of the others, and the aim is to select an agent nominated by many. Mechanisms are ranked according to approximation ratios[2] rather than according to expected values, and this leads to qualitatively different optimal mechanisms. For example, while jury mechanisms can be optimal in our model, the 2-*partition mechanism* of Alon et al. (2011), which is a natural analogue of jury mechanisms, is not optimal in their model.[3] [4]

See Olckers and Walsh (2022) for a survey of the literature following Holzman and Moulin (2013) and Alon et al. (2011). Olckers and Walsh also report on some related empirical studies.

Other work in mechanism design focuses on non-monetary instruments for eliciting information For example, in the aforementioned paper of Ben-Porath, Dekel, and Lipman (2014), the agents' types can be verified at a cost.[5] The typical assumption in this literature is that the agents do not have information about their peers. Most rel-

1. Further contributions to the literature following Holzman and Moulin (2013) include Tamura and Ohseto (2014), Tamura (2016), and Edelman and Por (2021). See also de Clippel, Moulin, and Tideman (2008).

2. Given $\alpha \in [0, 1]$, a mechanism has an *approximation ratio* of α if it guarantees a fraction α of some benchmark value. The guarantee is computed across all realizations of the type profile; that is, across all possible approval sets. The benchmark value at a particular realization is the maximal number of approvals across agents.

3. The 2-partition mechanism randomly splits the agents into two subsets, and then selects an agent from the first subset with the most approvals from agents in the second subset. Alon et al. (2011, Theorem 4.1) show that the 2-partition mechanism has an approximation ratio of $\frac{1}{4}$. Fischer and Klimm (2015) present a mechanism that achieves the strictly higher and optimal ratio of $\frac{1}{2}$.

4. Further contributions to this literature include Bousquet, Norin, and Vetta (2014), Aziz et al. (2016), Bjelde, Fischer, and Klimm (2017), Aziz et al. (2019), Mattei, Turrini, and Zhydkov (2020), and Lev et al. (2021). See also Caragiannis, Christodoulou, and Protopapas (2019, 2021), who consider additive approximations rather than approximation ratios.

5. See Epitropou and Vohra (2019), Erlanson and Kleiner (2019), and Li (2020) for further work with costly verification. Other examples of non-monetary instruments include promises of future allocations (Guo and Hörner, 2021), costly signaling (Condorelli, 2012; Chakravarty and Kaplan, 2013), allocative externalities (Bhaskar and Sadler, 2019; Goldlücke and Tröger, 2020), or ex-post punishments (Mylovanov and Zapechelnyuk, 2017; Li, 2020).

evant for us are papers that study how a *Bayesian* incentive-compatible mechanism may use agents' peer information to incentivize truthtelling (Kattwinkel, 2019; Kattwinkel and Knoepfle, 2021; Bloch, Dutta, and Dziubiński, 2022; Kattwinkel et al., 2022). The idea is that when agents have information about their peers, one can detect lies by cross-checking the agents' reports. We observe that the *dominant-strategy* incentive-compatible mechanisms that we consider do not use peer information in this manner. While DIC thus shuts down a screening channel, it leads to mechanisms that are far simpler for the agents to play. Relatedly, the fundamental insights of Crémer and McLean (1985, 1988) and McAfee and Reny (1992) on mechanisms *with transfers* do not apply here.

The papers of Baumann (2018) and Bloch and Olckers (2021, 2022) study related settings but focus on different questions. For instance, Bloch and Olckers (2022) study whether it is possible to reconstruct the ordinal ranking of agents from their reports when agents prefer a high rank.

We also contribute to the literature on the gap between stochastic and deterministic mechanisms[6] by fully characterizing when deterministic DIC mechanisms suffice for describing the set of DIC mechanisms in the present model. Methodologically, we show that here the existence of stochastic extreme points can be understood via a graph-theoretic result due to Chvátal (1975). We elaborate in Appendix 1.B.

1.3 Model

A single indivisible object is to be allocated to one of n agents, where $n \geq 2$. For each agent i, let Ω_i be a finite set of reals representing the possible social values from allocating to agent i, and let Θ_i be a finite set representing agent i's possible private types. Let $\Omega = \times_{i=1}^{n} \Omega_i$ and $\Theta = \times_{i=1}^{n} \Theta_i$. Values and types are distributed according to a joint distribution μ over $\Omega \times \Theta$. At all type profiles, agent i strictly prefers winning the object to not winning it; agent i is indifferent to which of the others is allocated the object.

In a *(direct) mechanism*, each agent reports a type, and then the object is allocated to one of the agents according to some lottery. Formally, a mechanism is a function $\phi : \Theta \to [0,1]^n$ satisfying $\sum_{i=1}^{n} \phi_i = 1$. Here $\phi_i : \Theta \to [0,1]$ denotes the winning probability of agent i. Since the object is allocated to one of the agents, these probabilities sum to 1. The requirement that the object is always allocated keeps with some earlier work (for example, Alon et al. (2011) and Holzman and Moulin (2013)). In Appendix 1.B, we discuss mechanisms that do not always allocate.

A mechanism ϕ is *dominant-strategy incentive-compatible (DIC)* if truthfully reporting one's type is a dominant strategy. For the assumed preferences of the agents,

6. See, for example, Budish et al. (2013), Pycia and Ünver (2015), Jarman and Meisner (2017), Chen et al. (2019), and Rivera Mora (2022).

a mechanism is DIC if and only if one's report never affects one's own winning probability.

To see the previous point in detail, let $u_i(\theta)$ denote the payoff to an agent i when i is allocated the object at a type profile θ. We normalize i's payoff when not allocated the object to 0, and we assume $u_i > 0$. DIC for a mechanism ϕ requires that all $i, \theta_i, \theta_i', \theta_{-i}$, and θ_{-i}' satisfy $u_i(\theta_i, \theta_{-i})\phi_i(\theta_i, \theta_{-i}') \geq u_i(\theta_i, \theta_{-i})\phi_i(\theta_i', \theta_{-i}')$. Since $u_i > 0$ and since θ_i and θ_i' are arbitrary, we must have $\phi_i(\theta_i, \theta_{-i}') = \phi_i(\theta_i', \theta_{-i}')$. That is, agent i's report never affects ϕ_i. Observe that nothing in this argument changes if $u_i < 0$. Hence we can equally model cases where some agents prefer not to be allocated the object.

We evaluate a DIC mechanism ϕ via the expected value of the allocation, which is given by $\mathbb{E}_{\omega,\theta}\left[\sum_{i=1}^{n} \phi_i(\theta)\omega_i\right]$. When we say a DIC mechanism is *optimal*, we mean it maximizes the expected value among all DIC mechanisms. The Revelation Principle implies that DIC mechanisms are without loss: if a mechanism can be implemented in some dominant-strategy equilibrium of some game, then it is DIC.

Lastly, we define the following: A mechanism is *deterministic* if it maps to a subset of $\{0, 1\}^n$. A mechanism is *stochastic* if it is not deterministic.

1.4 Jury mechanisms

In this section, we focus on the following class of mechanisms.

Definition 1.1. A mechanism ϕ is a *jury mechanism* if for all agents i we have the following: if the mechanism is non-constant in agent i's report, then agent i never wins, meaning $\phi_i = 0$.

Given a jury mechanism, we refer to an agent as a *juror* if the mechanism is non-constant in their report. The set of jurors is called the *jury*, and the remaining agents are called *candidates*. All jury mechanisms are DIC since jurors never win.

The most natural jury mechanisms are those that allocate to the top candidate conditional on the jurors' reports. That is, when the set of jurors is J and jurors report types $(\theta_i)_{i\in J}$, the object is allocated to one of the candidates in

$$\operatorname*{arg\,max}_{k\in\{1,\dots,n\}\setminus J} \mathbb{E}_{\omega_k}[\omega_k|(\theta_i)_{i\in J}].$$

Assuming a common prior, this mechanism would be implemented by having the jurors share their private information via cheap-talk messages, update their beliefs about the candidates, and then award the object to the top candidate given their shared posterior belief. (For our proofs, however, it is convenient to allow the jurors to select a suboptimal candidate.)

A priori, all agents in the model are candidates for winning and suppliers of information. Jury mechanisms are special since the roles of candidates and jurors are assigned *before* the agents are consulted. There are more complicated mechanisms where an agent's "role" varies across type profiles, and we shall encounter such mechanisms later. As such, it is remarkable that there are situations where jury mechanisms are (approximately) optimal, as we discuss next.

1.4.1 Jury mechanisms solve the three-agent case

Theorem 1.1. *Let $n \leq 3$. A mechanism is DIC if and only if it is a convex combination of deterministic jury mechanisms. In particular, there is an optimal DIC mechanism that is a deterministic jury mechanism.*

With three agents, a jury mechanism admits at most one juror who deliberates between the other two. Therefore, all DIC mechanisms with three agents can be implemented by nominating a juror (according to some distribution over the set of agents), and then asking the juror to pick one of the others as a winner of the object. Optimally, the information of at least two of the agents is ignored. (With only two agents, all DIC mechanisms are constant.)

In the remainder of this subsection, we explain the steps in the proof of Theorem 1.1. We begin with a known result (Holzman and Moulin, 2013, Proposition 2.i).

Lemma 1.2. *If $n \leq 3$, then all deterministic DIC mechanisms are jury mechanisms.*

In the language of Section 5 of Holzman and Moulin (2013), a deterministic DIC mechanism is an *impartial award rule*. Their Proposition 2.i implies that, if $n \leq 3$, then in each impartial award rule there is at most one agent whose report influences the allocation, and this influential agent never wins. Such a rule is a jury mechanism.[7]

To the best of our knowledge, Lemma 1.2 has so far been limited to deterministic DIC mechanisms. We now close the gap to stochastic ones.

Lemma 1.3. *If $n \leq 3$, then all DIC mechanisms are convex combinations of deterministic DIC mechanisms.*

Lemma 1.3 completes the proof of Theorem 1.1. Indeed, Lemma 1.2 and Lemma 1.3 immediately imply that all DIC mechanisms are convex combinations of deterministic jury mechanisms. Since the expected value is a linear function of the mechanism, at least one deterministic jury mechanism must be optimal.

To prove Lemma 1.3 we consider the extreme points of the set of DIC mechanisms. A routine argument shows that the set of DIC mechanisms is convex and

7. Holzman and Moulin (2013) note that the result is essentially due to Kato and Ohseto (2002), who study pure exchange economics. For a discussion of this relationship, we refer to Section 1.4 of Holzman and Moulin (2013).

compact (as a subset of Euclidean space). Hence, by the Krein-Milman theorem (Aliprantis and Border, 2006, Theorem 7.68), the set is given by the convex hull of its extreme points.

We show that all stochastic DIC mechanisms fail to be extreme points. Specifically, given an arbitrary stochastic DIC mechanism ϕ we construct a non-zero function f such that $\phi + f$ and $\phi - f$ are two other DIC mechanisms. To understand this construction, recall that a stochastic mechanism is one where, for at least one type profile, at least one agent enjoys an interior winning probability. Since the object is always allocated, some other agent must also enjoy an interior winning probability at the same profile. The function f represents a shift of a small probability mass between these two agents. This shift should be consistent with DIC (since we want $\phi + f$ and $\phi - f$ to be DIC), and hence we have to shift masses at multiple type profiles. What makes the construction of f difficult is that changing one agent's type may change which of the others enjoys an interior winning probability. Our argument thus intuitively leans on there only being three agents. Indeed, we shall later see that the argument does not go through with four or more agents.

1.4.2 Approximate optimality of jury mechanisms

In this subsection, we identify environments in which jury mechanisms are approximately optimal if the number n of agents is large. As suggested in the introduction, DIC creates a tension between allocating to an agent and using the agent's peer information. This tension becomes easier to resolve with many agents. Indeed, we intuit that many DIC mechanisms become approximately optimal as $n \to \infty$. The insight of the upcoming result is that this includes the DIC mechanisms that resolve the tension in the most straightforward way—jury mechanisms.

The following example conveys the basic idea.

Example 1.1. For each agent i, the value ω_i of allocating to i depends on some common component s and some private component t_i. Specifically, for some function $\hat{\omega}_i$ we have $\omega_i = \hat{\omega}_i(s, t_i)$ with probability 1. The agents observe their private components, which are independently and identically distributed across agents and independent of s. All agents observe s. (So, agent i's type is $\theta_i = (s, t_i)$.) Let ϕ be an arbitrary DIC mechanism for these n agents. Now suppose a new agent $n + 1$, who also observes the common component s, joins the group. Agent $n + 1$ may observe some additional information, but this will not be relevant. We claim that there is a jury mechanism that only uses agent $n + 1$ as a single juror and that does as well as ϕ. Note that, by ignoring the reports of agents 1 to n, the information contained in the public component s is not lost. The only information that is potentially lost is the first n agents' knowledge of their private components $t_1, \dots, t_n$. Each agent i's private component t_i is informative only about i's own value (by independence). However, DIC of the original mechanism ϕ implies that t_i could not have been used to determine i's own allocation. Thus one does not actually lose any information when ignoring the reports of agents 1 to n.

The main result of this section generalizes the previous example as follows. Under an assumption on the distribution of types and values, an arbitrary DIC mechanism with n agents can be replicated by a jury mechanism when additional agents are around. If values remain bounded in n, an implication is that the loss from using an optimal jury mechanism vanishes as $n \to \infty$.

We introduce new notation to accommodate the growing number of agents. The agents share a common finite type space ($\Theta_1 = \Theta_i$ for all i). The prior distribution of values and types is now a Borel-probability measure μ on $\times_{i\in\mathbb{N}}(\Omega_i \times \Theta_i)$,where each Ω_i is a finite set of reals.[8]

The following assumption captures the idea that if i, j, and k are three distinct agents, then i and j have access to the same sources of information about ω_k.

Assumption 1.1. For all $n \in \mathbb{N}$, all $i \in \{1, \ldots, n\}$, and all $\omega_i \in \Omega_i$, we have the following: Conditional on the value of agent i being equal to ω_i, the distribution of $(\theta_j)_{j\in\{1,\ldots,n\}\setminus\{i\}}$ is invariant with respect to permutations of $\{1, \ldots, n\} \setminus \{i\}$.

We are not assuming that i and j have the same information as k about ω_k. For example, in Example 1.1, the common component is the only information that i and j have about ω_k, but agent k actually observes ω_k.

When there are n agents (meaning that mechanisms only consult and allocate to the first n agents), let V_n denote the expected value from an optimal DIC mechanism. Let V_n^J denote the expected value from a jury mechanism with n agents that is optimal among jury mechanisms with n agents.

Theorem 1.4. *Let Assumption 1.1 hold. For all $n \in \mathbb{N}$ there exists $m \in \mathbb{N}$ such that $V_n \leq V_{n+m}^J$. If, additionally, the sequence $\{V_n\}_{n\in\mathbb{N}}$ is bounded,[9] then $\lim_{n\to\infty}(V_n - V_n^J) = 0$.*

In plain words, if m new agents are added to the group, a jury mechanism with $n + m$ agents does as well as an with an arbitrary DIC mechanism with n agents. The proof shows this claim for a jury mechanism that has the new m agents as jurors, and the old n agents as candidates, and where $m = n$. That is, a jury mechanism with the desired properties exists as soon as the number of agents is doubled. Depending on the exact distribution μ, a much smaller number of new agents may be needed; in Example 1.1, one new agent suffices.

Assumption 1.1 is stronger than what we really need. It suffices if, informally speaking, for all groups of agents $\{1, \ldots, n\}$ there eventually comes a disjoint group

8. Each of the finite sets Ω_i and Θ_i is equipped with the discrete metric. The product $\times_{i\in\mathbb{N}}(\Omega_i \times \Theta_i)$ is equipped with the product metric.

9. A sufficient condition for boundedness of the sequence $\{V_n\}_{n\in\mathbb{N}}$ is that the values ω_i are bounded across agents. For example, suppose with μ-probability 1 we have $\omega_i \in [0, 1]$ for all $i \in \mathbb{N}$.

of agents that is at least as well informed as $\{1,\ldots,n\}$ about each other. Assumption 1.2 in Appendix 1.A.1.2 formalizes this idea.

Remark 1. Theorem 1.4 does not assert that DIC mechanisms become approximately ex-post optimal conditional on the type profile. In Example 1.1, the only information that is used in the allocation is the common component. The common component need not pin down the entire profile of values.

1.5 Random allocations

In this section, we show that it typically does not suffice to consider deterministic mechanisms. This fact sheds light on the fundamental economic forces of the model and has practical implications for implementation, as we explain below.

1.5.1 Stochastic extreme points

One of way constructing a stochastic DIC mechanism is by randomizing over deterministic ones; that is, by taking a convex combination of deterministic DIC mechanisms. In this case, one of the deterministic mechanisms from the combination must generate a weakly higher expected value than the stochastic mechanism.

We therefore ask whether all stochastic DIC mechanisms can be represented as convex combinations of deterministic ones; that is, whether all extreme points of the set of DIC mechanisms are deterministic. In a nutshell, this is true if and only if there are at most three agents or the agents' type spaces are small.

Theorem 1.5. *All extreme points of the set of DIC mechanisms are deterministic if and only if at least one of the following is true:*

(1) There are at most three agents; that is, we have $n \leq 3$.

(2) All agents have at most two types; that is, for all i we have $|\Theta_i| \leq 2$.

(3) At least $(n-2)$-many agents have a degenerate type; that is, we have

$$|\{i \in \{1,\ldots,n\} : |\Theta_i| = 1\}| \geq n-2.$$

We already know from Lemma 1.3 that (1) is sufficient for all extreme points to be deterministic. Sufficiency of (2) is related to a generalization of the well-known Birkhoff-von Neumann theorem; sufficiency of (3) is economically and technically uninteresting, but must be included for completeness.[10] As for the other direction:

10. The reader may wonder whether one can prove sufficiency of (1) to (3) by viewing the set of DIC mechanisms as the set of solutions to a linear system of inequalities, checking for total unimodularity of the constraint matrix, and then invoking the Hoffman-Kruskal theorem (Korte and Vygen, 2018, Theorem 5.21). In the mechanism design literature, this approach is discussed in Pycia and Ünver (2015), for example. Here the approach works for the case where all type spaces are binary; our proof uses a result which can itself be derived from the Hoffman-Kruskal theorem. However, in the difficult case with three agents, the constraint matrix is not generally totally unimodular (see Appendix 1.C.3).

we momentarily give an example of a stochastic extreme point. The general claim that a stochastic extreme point exists when (1) to (3) all fail follows readily by extending this example.

An implication of Theorem 1.5 is that deterministic DIC mechanisms do not suffice for optimality. Indeed, for each extreme point there exists at least one distribution of types and values where the extreme point is the unique optimal DIC mechanisms.[11]

We do not expect stochastic extreme points to closely resemble mechanisms observed in practice. The literature discusses several issues. First, to reduce complexity and opaqueness, it is appealing to implement a mechanism by randomizing over deterministic mechanisms, announcing the selected mechanism, and only then collecting the agents' reports (see, for example, Pycia and Ünver (2015)). A stochastic extreme point is precisely a DIC mechanism that cannot be implemented in this way.[12] Second, to implement a stochastic extreme point, the designer must commit to honoring the outcome of a stochastic process (see, for example, Chen et al. (2019)). A commitment issue arises if the agents' collective information identifies a unique qualified agent but the mechanism nevertheless promises to flip a coin between this agent and a less qualified one.

Despite the above points, it may be acceptable to randomize if this happens "rarely" or is used to break ties between "similar" agents. As it happens, the optimality of stochastic extreme points is not limited to such cases. We next present an example where a stochastic extreme point is uniquely optimal. This stochastic extreme point "frequently" randomizes between "dissimilar" agents.

1.5.2 An example of a stochastic extreme point

There are four agents, and their types are as follows:

$$\Theta_1 = \{\ell, r\}, \quad \Theta_2 = \{u, d\}, \quad \Theta_3 = \{f, c, b\}, \quad \Theta_4 = \{0\}. \tag{1.1}$$

Figure 1.1 shows (among other things that are not yet relevant) the type profiles of agents 1, 2, and 3; the degenerate type of agent 4 is omitted. The types of agents

11. The argument is as follows. The set of DIC mechanisms is a polytope in Euclidean space that does not depend on the distribution. All extreme points of the polytope are exposed. Since all linear functionals on this polytope can be represented via some distribution, the claim follows. See Appendix 1.C.1 for the formalities.

12. In fact, in our model, stochastic extreme points cannot be implemented via any dominant-strategy equilibrium of any deterministic *indirect* mechanism. See Appendix 1.C.2. We note, however, a result of Rivera Mora (2022) implying the following (for our model): *Given an arbitrary DIC direct mechanism, there is an* ex-post *equilibrium of a deterministic indirect mechanism that implements the given DIC direct mechanism.* In this ex-post equilibrium, the agents play mixed strategies that emulate the randomization on the part of the given DIC mechanism. These mixed strategies do not generally form a dominant-strategy equilibrium.

1, 2, and 3 span a three-dimensional hyperrectangle. (Mnemonically, their types mean left, right, up, down, front, center, and back.) Each edge of the hyperrectangle represents a set of type profiles along which exactly one agent's type is changing. Hence DIC requires that the winning probability of this agent be constant along the edge. We identify such an edge by a pair (i, θ_{-i}), where i indicates the agent whose type is changing, and θ_{-i} indicates the fixed types of the others.

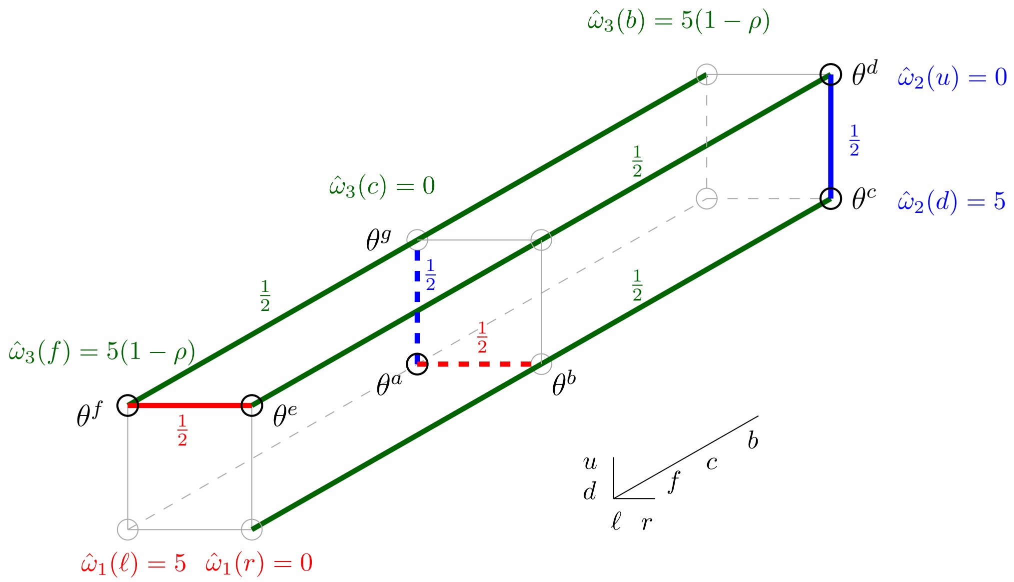

Figure 1.1. The set of types of agents 1, 2, and 3. The probabilities $\frac{1}{2}$ attached to the edges of the hyperrectangle represent the relevant values of the mechanism ϕ^*. The values from the allocation are as defined in (1.5). The distribution μ assigns probability $\frac{1}{5}$ to the profiles $\{\theta^a, \theta^c, \theta^d, \theta^e, \theta^f\}$. All other profiles have probability 0.

Let $\Theta^* = \{\theta^a, \theta^b, \theta^c, \theta^d, \theta^e, \theta^f, \theta^g\}$ be the set of labeled type profiles in Figure 1.1; these are the profiles

$$\begin{aligned} \theta^a &= (\ell, d, c, 0), & \theta^b &= (r, d, c, 0), & \theta^c &= (r, d, b, 0), \\ \theta^d &= (r, u, b, 0), & \theta^e &= (r, u, f, 0), & \theta^f &= (\ell, u, f, 0), \\ \theta^g &= (\ell, u, c, 0). \end{aligned} \tag{1.2}$$

Let V^* denote the set of bold edges in Figure 1.1 that connect the profiles in Θ^*; these are the edges

$$V^* = \{(1, \theta^a_{-1}), (3, \theta^c_{-3}), (2, \theta^c_{-2}), (3, \theta^e_{-3}), (1, \theta^e_{-1}), (3, \theta^f_{-3}), (2, \theta^a_{-2})\}.$$

Our candidate stochastic extreme point ϕ^* is defined as follows (see Figure 1.1): For all $i \in \{1, 2, 3\}$ and $\theta \in \Theta$, let

$$\phi^*_i(\theta) = \begin{cases} \frac{1}{2}, & \text{if } (i, \theta_{-i}) \in V^*, \\ 0, & \text{otherwise.} \end{cases}$$

Further, for all $\theta \in \Theta$ let $\phi_4^*(\theta) = 1 - \sum_{i \in \{1,2,3\}} \phi_i^*(\theta)$. In plain words, at all profiles in Θ^*, exactly two bold edges of the hyperrectangle intersect at the profile; the mechanism ϕ^* randomizes evenly between the two agents of these edges. All remaining probability mass is assigned to agent 4. It is easy to verify from Figure 1.1 that ϕ^* is a well-defined DIC mechanism.

Further below we specify values Ω and a distribution μ such that ϕ^* is the unique optimal DIC mechanism. This implies that ϕ^* is an extreme point of the set of DIC mechanisms. Since the proof for uniqueness is somewhat involved, we next present a simple self-contained argument showing that ϕ^* is an extreme point.

Let ϕ be a DIC mechanism that receives non-zero weight in a convex combination that equals ϕ^*. We show $\phi = \phi^*$. For all profiles $\theta \in \Theta^*$, there are exactly two agents i and j such that (i, θ_{-i}) and (j, θ_{-j}) both belong to V^*; these are the two bold edges of the hyperrectangle that intersect at θ. Hence at θ the mechanism ϕ^* randomizes evenly between i and j. Since ϕ is part of a convex combination that equals ϕ^*, it follows that at θ the mechanism ϕ only randomizes between i and j, meaning $\phi_i(\theta) = 1 - \phi_j(\theta)$. Since ϕ is DIC, repeatedly applying this observation shows:

$$\begin{aligned}
\phi_1(\theta^a) \;=\; 1 - \phi_3(\theta^c) \;=\; \phi_2(\theta^c) \;&=\; 1 - \phi_3(\theta^e) \\
&=\; \phi_1(\theta^e) \\
&=\; 1 - \phi_3(\theta^f) \;=\; \phi_2(\theta^a) \;=\; 1 - \phi_1(\theta^a).
\end{aligned} \tag{1.3}$$

In particular, we have $\phi_1(\theta^a) = 1 - \phi_1(\theta^a)$, implying $\phi_1(\theta^a) = \frac{1}{2}$. Hence all probabilities in (1.3) equal $\frac{1}{2}$. Hence ϕ agrees with ϕ^* at all profiles in Θ^*. By inspecting $\Theta \setminus \Theta^*$, we may easily convince ourselves that ϕ and ϕ^* also agree on $\Theta \setminus \Theta^*$. Thus ϕ^* is an extreme point.

We next construct an environment in which ϕ^* is uniquely optimal. We could do so by invoking a separating hyperplane theorem. However, this would be unsatisfying since we would gain no intuition for why randomization helps or for whether ϕ^* is uniquely optimal in a restricted class of environments. We shall gain both by considering environments in which values are *privately known*, in the following sense: for all agents i, the value of allocating to i is pinned down by a function $\hat{\omega}_i$ that depends only on θ_i.

We can describe an environment with privately known values by specifying a distribution μ over type profiles and, for all agents i, a function $\hat{\omega}_i \colon \Theta_i \to \mathbb{R}$ that governs the value of allocating to i. Our candidate distribution μ is given by (see Figure 1.1)

$$\forall_{\theta \in \Theta}, \quad \mu(\theta) \;=\; \begin{cases} \frac{1}{5}, & \text{if } \theta \;\in\; \{\theta^a, \theta^c, \theta^d, \theta^e, \theta^f\} \\ 0, & \text{else.} \end{cases} \tag{1.4}$$

Our candidates for $\hat{\omega}_1, \ldots \hat{\omega}_4$ are parametrized by $\rho \in [0, \frac{1}{2}]$ and given by

$$\begin{aligned} \hat{\omega}_1(r) = \hat{\omega}_2(u) = \hat{\omega}_3(c) &= 0 \\ \hat{\omega}_1(\ell) = \hat{\omega}_2(d) &= 5 \\ \hat{\omega}_3(f) = \hat{\omega}_3(b) &= 5(1-\rho) \\ \hat{\omega}_4 &= 0. \end{aligned} \tag{1.5}$$

Proposition 1.6. *The mechanism ϕ^* is an optimal DIC mechanism if and only if $\rho \in [0, \frac{1}{2}]$, and it is uniquely optimal if and only if $\rho \in (0, \frac{1}{2})$.*

In the introduction, we intuited that there is a trade-off between allocating to an agent and using that agent's information about others. In the present example, this trade-off involves agent 3 and depends on ρ.

To gain an intuition for the trade-off and the result, consider the case $\rho = 0$. Allocating to agent 3 is now ex-post optimal at *all except one* of the five profiles in the support of μ. Indeed, one optimal DIC mechanisms is the constant one that always allocates to agent 3. The mechanism ϕ^* is another optimal mechanism for $\rho = 0$, which is intuitively explained by agent 3's type being informative: if $\theta_3 = c$ realizes, the type profile must be θ^a, where θ^a is the unique type profile in the support of μ at which allocating to agents 1 or 2 is better than allocating to agent 3. The mechanism ϕ^* indeed allocates to agents 1 and 2 at θ^a.

Since ρ decreases the value from allocating to agent 3, it is now intuitive that ϕ^* does strictly better than always allocating to agent 3 for small but strictly positive values of ρ. In the formal proof, most of our effort goes towards showing that ϕ^* is in fact uniquely optimal for small but strictly positive values of ρ. The idea is that, among all DIC mechanisms that are optimal for $\rho = 0$, the mechanism ϕ^* is the unique one minimizing agent 3's overall winning probability.

If we increase ρ further, it eventually becomes optimal to use agent 3 as a source of information and never allocate to agent 3. The critical value turns out to be $\rho = \frac{1}{2}$. The intuition is confirmed by the fact that, if $\rho = \frac{1}{2}$, the following *jury* mechanism with agent 3 as a juror is optimal: if agent 3 reports f, agent 1 wins; if agent 3 reports c, a coin flip determines whether agent 1 or 2 wins; if agent 3 reports b, agent 2 wins.

Proposition 1.6 also helps illustrate the commitment issue discussed in the paragraphs following Theorem 1.5. At the profile θ^e, a coin flip determines whether agent 1 or 3 wins the object. Yet, at this profile, the value from allocating to agent 3 is strictly higher than the value from allocating to agent 1. In fact, a coin is flipped at all type profiles in the support of the distribution. For $\rho \in (0, \frac{1}{2})$, the mechanism designer is indifferent to the outcome of the coin flip at only one of these profiles.

Remark 2. Chen et al. (2019) show that, in certain mechanism design problems, given any stochastic mechanism there is a deterministic one that induces the same interim-expected allocations. Since the deterministic mechanism is not guaranteed to be DIC, their result does not contradict the suboptimality of deterministic DIC mechanisms in our model.

Remark 3. An alternative approach to showing the existence of a stochastic extreme point uses a graph-theoretic result due to Chvátal (1975), as we explain in Appendix 1.B.2. For a certain graph G that we define in Appendix 1.B.2, Chvátal's theorem implies that all extreme points are deterministic if and only if G is perfect. To be precise, the results of this appendix concern the related problem where the mechanism may dispose the object instead of allocating it to the agents. The associated characterization of extreme points is implied by Theorem 1.5, but not vice versa.

1.6 Anonymous juries

In this section, we study anonymous DIC mechanisms. Anonymity, formally defined below, is roughly the requirement that any two agents exert the same influence with their reports on the winning probability of any third agent. This is a desirable property as it helps protect the agents' privacy when they evaluate their peers, reduces the complexity of the mechanism, and ensures that the agents have the same voting rights.

We offer two insights. First, all anonymous DIC mechanisms ignore the reports of the agents. Second, we consider a relaxed notion of anonymity—*partial anonymity*—and show that all deterministic partially anonymous DIC mechanisms are jury mechanisms.

Throughout, we assume that the agents share a common type space, meaning $\Theta_1 = \ldots = \Theta_n$. In an equally valid interpretation, we can consider indirect mechanisms where all agents have the same message space and cannot influence their own winning probabilities.

1.6.1 Notions of anonymity

Anonymity and partial anonymity are defined next. Anonymity requires that, for all k, the winning probability of agent k does not change if one permutes the reports of the agents other than k. Partial anonymity relaxes anonymity as follows: When testing whether k's winning probability is affected by permutations, we only consider permutations of those agents who actually influence agent k. In particular, partial anonymity permits the set of agents who influence k to be a proper subset of $\{1, \ldots, n\} \setminus \{k\}$.

Definition 1.2. Let the agents have a common type space. Let ϕ be a mechanism.

(1) Given i, j, and k that are all distinct, agents i and j are *exchangeable for k* if ϕ_k is invariant with respect to permutations of i's and j's reports; that is, for all profiles

θ and θ' such that θ is obtained from θ' by permuting the types of i and j we have $\phi_k(\theta) = \phi_k(\theta')$.

(2) Given distinct i and k, agent i *influences* k if ϕ_k is non-constant in i's report; that is, there exist type profiles θ and θ' that differ only in i's type and satisfy $\phi_k(\theta) \neq \phi_k(\theta')$.

(3) The mechanism is *anonymous* if for all i, j, and k that are all distinct, agents i and j are exchangeable for k.

(4) The mechanism is *partially anonymous* if for all i, j, and k that are all distinct we have the following: if i and j both influence k, then i and j are exchangeable for k.

To state the upcoming characterization of partial anonymity, we also define what we mean by an anonymous jury.

Definition 1.3. Let the agents have a common type space. A jury mechanism has an *anonymous jury* if all jurors i and j are exchangeable for all agents k.

Remark 4. If Assumption 1.1 holds, then among jury mechanisms it is without loss to use one with an anonymous jury. Indeed, consider the jury mechanism that selects the candidate that is best conditional on the types of the jurors (breaking ties in some fixed order). Under Assumption 1.1, the identity of the favored candidate does not change when one permutes the jurors' types.

1.6.2 Anonymous DIC mechanisms ignore all reports

Theorem 1.7. *Let the agents have a common type space. All anonymous DIC mechanisms are constant.*

Note well that anonymity does *not* demand that i and j be exchangeable for i's own winning probability. If anonymity did demand this, the theorem would follow rather trivially from DIC.

The theorem is more subtly related to the requirement that the mechanism always allocates the object, as we explain next. This requirement lets us link the influence that two agents i and j exert on others to the influence that they exert on each other.

More concretely, assume towards a contradiction that at some profile θ agent i can increase ϕ_j by changing their report from θ_i to some θ_i'. By DIC and since the object is always allocated, this change in i's report decreases $\sum_{k:\, i \neq k \neq j} \phi_k$. Now consider the profile that is obtained from θ by permuting the reports of i and j. By anonymity, agent j can change their report from θ_i to θ_i' to decrease $\sum_{k:\, i \neq k \neq j} \phi_k$. Using again that the mechanism is DIC and that the object must be allocated, it follows that the change in agent j's report increases ϕ_i. In summary, if i can increase j's winning probability at some profile, then j must also be able to increase i's winning probability at a permuted profile. This observation suggests that i and j both win with “high” probability when both report θ_i'. In a deterministic mechanism, where winning probabilities are either 0 or 1, we thus arrive at a contradiction to there being only one object to allocate. We address stochastic mechanisms via a substantially more complex summation over winning probabilities across all pairs (i,j).

Remark 5. Theorem 1.7 implies that all DIC mechanisms satisfying the following stronger notion of anonymity are constant: Whenever the set of reports is permuted, then the same permutation is applied to the vector of winning probabilities. This stronger notion captures a sense in which agents are treated equally both as voters and winners.

Remark 6. An implication of Theorem 1.7 is that it is impossible to elicit information in environments where anonymity is without loss. Indeed, if the joint distribution of types and values is invariant with respect to all permutations of the agents, then it is without loss to use a DIC mechanism that satisfies the strong notion of anonymity from Remark 5. Hence in this case it is without loss to use a constant mechanism.

1.6.3 Partial anonymity and jury mechanisms

Theorem 1.7 implies that a non-constant DIC mechanism must admit some asymmetry in how it processes the reports of different agents. This brings us to partial anonymity. We offer the following characterization for deterministic mechanisms.

Theorem 1.8. *Let the agents have a common type space. A mechanism is deterministic, partially anonymous, and DIC if and only if it is a deterministic jury mechanism with an anonymous jury.*

To better understand the theorem, consider how a partially anonymous jury mechanism could fail to admit an anonymous jury. Given agents i and j, partial anonymity is silent on the winning probabilities of those agents k who are influenced by either i or j but not by both. By contrast, anonymity of the jury requires that all candidates are either influenced by all or none of the jurors. Accordingly, most of our effort goes towards proving that, in a deterministic partially anonymous DIC mechanism, if i and j influence *some* third agent k, then i and j influence exactly the same set of agents. Equipped with this fact, we show that the agents can be partitioned into equivalence classes with the following property: two agents in the same class do not influence one another, but influence the same (possibly empty) set of agents outside the class. Lastly, there cannot be multiple classes; indeed, else there is a profile where two distinct classes allocate the object to two distinct agents, which is impossible. The unique class defines an anonymous jury.

1.6.4 Discussion and limitations

We conclude by discussing limitations of Theorem 1.7 and Theorem 1.8.

1.6.4.1 Disposal and randomization

The following definition will be useful: A *mechanism with disposal* is a function $\phi : \Theta \to [0,1]^n$ satisfying $\sum_{i=1}^n \phi_i \leq 1$. In plain words, this is a mechanism that does not necessarily allocate the object to the agents. For a mechanism with disposal, DIC and anonymity are defined as above.

The next result shows via an example that Theorem 1.7 does not extend to mechanisms with disposal, and that Theorem 1.8 does not extend to stochastic mechanisms (without disposal).

Proposition 1.9. *Let the agents have a common type space T such that $|T| = 7$.*

(1) If $n = 3$, then the set of DIC mechanisms with disposal admits an extreme point that is stochastic and anonymous.
(2) If $n = 4$, then the set of DIC mechanisms (without disposal) admits an extreme point that is stochastic and partially anonymous.

The extreme point in (1) is non-constant (else it would be a convex combination of deterministic constant mechanisms). The extreme point in (2) is not a jury mechanism (else it would be a convex combination of deterministic jury mechanisms). The idea of the proof is to "symmetrize" the stochastic extreme point ϕ^* from Section 1.5.2. See Appendix 1.A.3.3 for an informal sketch and the proof.

1.6.4.2 Anonymous ballots

Lastly, we discuss the assumption that all agents can make the same reports. Indeed, a third escape route from Theorem 1.7 (besides partial anonymity and disposal) entails message spaces with some inherent asymmetry across agents. This brings us to the results of Holzman and Moulin (2013) and Mackenzie (2015, 2020). They consider DIC mechanisms where agents nominate one another. Let us keep with the terminology of Holzman and Moulin by referring to these mechanisms as *impartial nomination rules*. This is the same mathematical object as a DIC mechanism when each agent i's type space is $\{1, \dots, n\} \setminus \{i\}$. Their notion of anonymity—*anonymous ballots*—requires that the winning probabilities depend only on the number of nominations received by each agent.[13] Importantly, in a nomination rule agents cannot nominate themselves, and hence they all have distinct message spaces. By contrast, we have assumed that the agents have the same type space. Hence our notion of anonymity neither nests nor is nested by anonymous ballots.

Contrasting Theorem 1.7, there are *non-constant* impartial nomination rules with anonymous ballots. For one example, suppose one of the agents is selected

13. Equivalently, the allocation is unchanged if one permutes the profile in a way that does not yield self-nominations (Mackenzie, 2015, Lemma 1.1). Mackenzie uses the name *voter anonymity* instead of anonymous ballots.

uniformly at random as a juror, following which the juror's nomination determines a winner. See Mackenzie (2015, Theorem 1) for a full characterization of anonymous ballots. Mackenzie's result generalizes Theorem 3 of Holzman and Moulin (2013), who had previously shown that all *deterministic* impartial nomination rules with anonymous ballots are constant.

Mackenzie (2020) shows that impartiality and anonymous ballots are compatible for deterministic nomination rules with disposal.[14] This parallels our discussion from Section 1.6.4 and contrasts the aforementioned Theorem 3 of Holzman and Moulin (2013). Mackenzie (2020, Theorem 1) also shows that when agents can nominate themselves, then deterministic impartial nomination rules with anonymous ballots must be constant. This is a special case of our Theorem 1.7 as anonymous ballots with self-nominations is stronger than anonymity.

1.7 Conclusion

We saw that jury mechanisms are optimal with three agents, and approximately-optimal when there are many exchangeable agents in the sense of Assumption 1.1. While DIC mechanisms cannot process all reports anonymously, jury mechanisms are the only deterministic partially anonymous DIC mechanisms. Lastly, outside of special cases of the model, the set of DIC mechanisms admits stochastic extreme points.

We conclude by discussing some interesting open problems.

The discussion on stochastic extreme points (Section 1.5.1) motivates restricting attention to deterministic mechanisms. We observe in Appendix 1.C.4 that finding an optimal deterministic DIC mechanism can be cast as the problem of finding a maximum weight perfect matching in a certain hypergraph. If we relax the requirement that the object is always allocated, the problem can also be cast as finding a maximum weight independent set in another graph. Both of these problems are known to be NP-hard when general (hyper-)graphs and weights are considered. As such, it is interesting to investigate the hardness of the problem for the particular family of (hyper-)graphs that emerge from our model. (All weights can emerge via a suitable choice of the distribution of types and values.) If we include stochastic mechanisms in our search, finding an optimal DIC mechanism is a linear program and hence computationally tractable.

It is naturally interesting to extend the analysis to settings with multiple objects, allocated simultaneously or over many periods.[15] If the mechanism designer

14. In fact, Mackenzie (2020, Theorem 2) shows that impartiality, anonymous ballots, and some other desirable axioms together characterize supermajority.

15. See Guo and Hörner (2021) for recent work in this direction with a single agent. The literature following Alon et al. (2011) has also studied settings with multiple objects. Lipnowski and Ramos (2020) and de Clippel et al. (2021) study settings with limited or no commitment.

can commit to future allocations, this should lead to stronger foundations for jury mechanisms. Agents serving as jurors today can be promised a future spot as candidates, which may help justify excluding jurors as potential winners in the present. Alternatively, past winners may be expected to volunteer as jurors in the future.

The problem of finding an optimal composition of the jury is an interesting problem in itself. We expect interesting comparative statics when agents who are likely to have good information are also likely to yield a high value. In the example from the introduction where a group selects a president, say, an agent who is popular with others may be a suitable candidate (being well-liked for their pleasant qualities) but also have good information about others (being well-acquainted with everyone).

An important line of future research concerns optimal DIC mechanisms when agents care about the allocation to their peers. While DIC has different implications in such a model, our results provide insight in at least two cases. Firstly, in situations where agents evaluate their peers, it is seems inherently interesting to use a mechanism where agents cannot influence their individual chances of winning; that is, to impose the *impartiality* axiom of Holzman and Moulin (2013). Secondly, suppose agents have the following lexicographic preferences: each agent i strictly prefers one allocation to another if the former has i winning with strictly higher probability; if two allocations have the same winning probability for i, agent i ranks them according to some type-dependent preference. In some applications, this preference could reasonably capture i's opinion about who is the most deserving winner if it cannot be i themself. In particular, it could coincide with the preference of the mechanism designer. In this case, optimal jury mechanism are ex-post incentive compatible. However, an agent's preferences may also differ from those of the designer. This is plausibly the case when agents are biased in favor of friends or family, biased against minorities, or simply have a different notion of who deserves to win.[16] Fixing a jury of agents, the designer therefore also has to design a voting rule for eliciting the jurors' information.

16. For example, Alatas et al. (2012), reporting on a field experiment on selecting beneficiaries of aid programs in Indonesian communities, find evidence of nepotism, though the welfare impact may be small relative to other upsides from involving the community in the decision. They also find evidence that community members have a poverty notion that differs from poverty as defined by per capita income. In this sense, if the central government wishes to select beneficiaries on the basis of per capita income, agents indeed hold a different notion of who deserves to win.

Appendix 1.A Omitted proofs

1.A.1 Jury mechanisms

1.A.1.1 Proof of Lemma 1.3

Proof of Lemma 1.3. If $n = 1$ or $n = 2$, it is easy to verify that all DIC mechanisms are constant. All constant mechanisms are convex combination of deterministic constant mechanisms, proving the claim. In what follows, let $n = 3$. Given an arbitrary stochastic DIC mechanism ϕ, we will find a non-zero function f such that $\phi + f$ and $\phi - f$ are two other DIC mechanisms. This shows that all extreme points of the set of DIC mechanisms are deterministic. Since this set is non-empty, convex and compact as a subset of Euclidean space, the claim follows from the Krein-Milman theorem.

In what follows, we fix a stochastic DIC mechanisms ϕ. Let us agree to the following terminology. In view of DIC, we drop i's type from ϕ_i. Given a profile θ, we refer to the equation $\sum_{i\in\{1,2,3\}} \phi_i(\theta_{-i}) = 1$ as the *feasibility constraint* at profile θ. We refer to (i, θ_{-i}) as the *node of agent i with coordinates θ_{-i}*. Lastly, when we say $\phi_i(\theta_{-i})$ is interior we naturally mean $\phi_i(\theta_{-i}) \in (0, 1)$.

Most of the work will go towards proving the following auxiliary claim.

Claim 1.10. *There are non-empty disjoint subsets R and B ("red" and "blue") of $\cup_{i\in\{1,2,3\}}(\{i\} \times \Theta_{-i})$ such that all of the following are true:*

(1) If $(i, \theta_{-i}) \in R \cup B$, then $\phi_i(\theta_{-i})$ is interior.

(2) For all $\theta \in \Theta$, exactly one of the following is true:

a. There does not exist $i \in \{1, 2, 3\}$ such that $(i, \theta_{-i}) \in R \cup B$.

b. There exists exactly one $i \in \{1, 2, 3\}$ such that $(i, \theta_{-i}) \in R$, exactly one $j \in \{1, 2, 3\}$ such that $(j, \theta_{-j}) \in B$, and exactly one $k \in \{1, 2, 3\}$ such that $(k, \theta_{-k}) \notin R \cup B$.

Before proving Claim 1.10, let us use it to complete the proof of Lemma 1.3. For a number ε to be chosen in a moment, let $f : \Theta \to \{-\varepsilon, 0, \varepsilon\}^3$ be defined as follows:

$$\forall_{\theta\in\Theta}, \quad f_i(\theta) = \begin{cases} -\varepsilon, & \text{if } (i, \theta_{-i}) \in R, \\ \varepsilon, & \text{if } (i, \theta_{-i}) \in B, \\ 0, & \text{if } (i, \theta_{-i}) \notin R \cup B. \end{cases}$$

By finiteness of Θ and Claim 1.10, if we choose $\varepsilon > 0$ sufficiently close to 0, then $\phi + f$ and $\phi - f$ are two DIC mechanisms. Since f is non-zero, it follows that ϕ is not an extreme point. It remains to prove Claim 1.10.

Proof of Claim 1.10. Given candidate sets R and B, let us say a profile θ is *uncolored* if it falls into case (2.a) of Claim 1.10. A profile *two-colored* if it falls into case (2.a) of Claim 1.10. In this terminology, our goal is to construct sets R and B such that all $(i, \theta_{-i}) \in R \cup B$ satisfy $\phi_i(\theta_{-i}) \in (0,1)$, and such that all type profiles are either uncolored or two-colored.

Since ϕ is stochastic, we may assume (after possibly relabelling the agents and types) that there exists a profile θ^0 such that $\phi_1(\theta_2^0, \theta_3^0)$ and $\phi_2(\theta_1^0, \theta_3^0)$ are interior.

Let Θ_2° denote the set of types θ_2 for which $\phi_1(\theta_2, \theta_3^0)$ is interior. Let $\Theta_2^\partial = \Theta_2 \setminus \Theta_2^\circ$. Similarly, let Θ_1° denote the set of types θ_1 such that $\phi_2(\theta_1, \theta_3^0)$ is interior, and let $\Theta_1^\partial = \Theta_1 \setminus \Theta_1^\circ$. Notice that Θ_1° and Θ_2° are non-empty as, by assumption, agents 1 and 2 are enjoying interior winning probabilities at θ^0.

We consider two cases.

Case 1. Let $\Theta_1^\partial \neq \emptyset$ and $\Theta_2^\partial \neq \emptyset$.

We establish two auxiliary claims.

Claim 1.11. *If $\theta_1 \in \Theta_1^\partial$, then $\phi_2(\theta_1, \theta_3^0) = 0$. Similarly, if $\theta_2 \in \Theta_2^\partial$, then $\phi_1(\theta_2, \theta_3^0) = 0$. If $(\theta_1, \theta_2) \in (\Theta_1^\circ \times \Theta_2^\partial) \cup (\Theta_1^\partial \times \Theta_1^\circ)$, then $\phi_3(\theta_1, \theta_2)$ is interior.*

Proof of Claim 1.11. Consider the first part of the claim. Let $\theta_1 \in \Theta_1^\partial$. Recalling that Θ_1° is non-empty, let us find a type $\theta_2 \in \Theta_1^\circ$. By definition, $\phi_1(\theta_2, \theta_3^0)$ is interior. By definition of Θ_1^∂, we also know that $\phi_2(\theta_1, \theta_3^0)$ must either equal 0 or 1. But it cannot equal 1 since $\phi_2(\theta_1, \theta_3^0)$ and $\phi_1(\theta_2, \theta_3^0)$ both appear in the feasibility constraint at the profile $(\theta_1, \theta_2, \theta_3^0)$, and since $\phi_1(\theta_2, \theta_3^0)$ is interior. Thus $\phi_2(\theta_1, \theta_3^0) = 0$, as desired.

A similar argument establishes the second claim.

As for the third claim, let $(\theta_1, \theta_2) \in \Theta_1^\circ \times \Theta_2^\partial$. The previous two paragraphs imply that at the profile $(\theta_1, \theta_2, \theta_3^0)$ the winning probability of agent 1 is 0. Moreover, by definition of Θ_1°, the winning probabiltiy of agent 2 is interior. Thus agent 3's winning probability at this profile must be interior, meaning $\phi_3(\theta_1, \theta_2)$ is interior. A similar argument shows that $\phi_3(\theta_1, \theta_2)$ is interior whenever (θ_1, θ_2) is in $\Theta_1^\partial \times \Theta_1^\circ$. □

The second auxiliary result is:

Claim 1.12. *Let $\theta_3 \in \Theta_3$. If $\theta_2 \in \Theta_2^\circ$, then $\phi_1(\theta_2, \theta_3)$ is interior. Similarly, if $\theta_1 \in \Theta_1^\circ$, then $\phi_2(\theta_1, \theta_3)$ is interior.*

Proof of Claim 1.12. We will prove the first part of the claim, the second being similar. Thus let $\theta_2 \in \Theta_2^\circ$. By assumption of Case 1, we may find $\theta_1^\partial \in \Theta_1^\partial$ and $\theta_2^\partial \in \Theta_2^\partial$. We make two auxiliary observations.

First, consider the profile $(\theta_1^\partial, \theta_2^\partial, \theta_3^0)$. According to Claim 1.11, both agent 1's and agent 2's winning probabilities at this profile equal 0. Thus $\phi_3(\theta_1^\partial, \theta_2^\partial) = 1$. But $\phi_3(\theta_1^\partial, \theta_2^\partial)$ and $\phi_2(\theta_1^\partial, \theta_3)$ both appear in the feasibility constraint at the profile $(\theta_1^\partial, \theta_2^\partial, \theta_3)$. Hence $\phi_2(\theta_1^\partial, \theta_3) = 0$.

Second, since $\theta_1^\partial \in \Theta_1^\partial$ and $\theta_2 \in \Theta_2^\circ$, we infer from Claim 1.11 that $\phi_3(\theta_1^\partial, \theta_2)$ is interior.

The previous two observations imply that at the profile $(\theta_1^\partial, \theta_2, \theta_3)$ agent 2's winning probability is 0 and that agent 3's winning probability is interior. Hence $\phi_1(\theta_2, \theta_3)$ is interior, as promised. □

We are ready to define the sets R and B. We assign the following colors (recall the terminology introduced in the paragraph before Claim 1.10):

- red to all nodes of agent 1 with coordinates in $\Theta_2^\circ \times \Theta_3$,
- blue to all nodes of agent 3 with coordinates in $\Theta_1^\partial \times \Theta_2^\circ$,
- blue to all nodes of agent 2 with coordinates in $\Theta_1^\circ \times \Theta_3$,
- red to all nodes of agent 3 with coordinates in $\Theta_1^\circ \times \Theta_2^\partial$.

According to Claim 1.11 and Claim 1.12, all of these nodes are interior. Moreover, all profiles are now either two-colored or uncolored: The profiles in $\Theta_1^\partial \times \Theta_2^\circ \times \Theta_3$ are two-colored via red nodes of agent 1 and blue nodes of agent 3; the profiles in $\Theta_1^\circ \times \Theta_2^\circ \times \Theta_3$ are two-colored via red nodes of agent 1 and blue nodes of agent 2; the profiles in $\Theta_1^\circ \times \Theta_2^\partial \times \Theta_3$ are two-colored via blue nodes of agent 2 and red nodes of 3; and the profiles in $\Theta_1^\partial \times \Theta_2^\partial \times \Theta_3$ are uncolored.

Case 2. Suppose at least one of the sets Θ_1^∂ and Θ_2^∂ is empty. In what follows, we assume that Θ_2^∂ is empty, the other case being analogous (switch the roles of agents 1 and 2).

The assumption that Θ_2^∂ is empty means that $\phi_1(\theta_2, \theta_3^0)$ is interior for all θ_2. Let Θ_1^* be the set of types θ_1 such that for all $\theta_2 \in \Theta_2$ the probability $\phi_3(\theta_1, \theta_2)$ is interior. Notice that at this point Θ_1^* may or may not be empty; we will make a case distinction further below.

We first claim that if $\theta_1 \in (\Theta_1 \setminus \Theta_1^*)$, then $\phi_2(\theta_1, \theta_3^0)$ is interior. Towards a contradiction, suppose this were false for some $\theta_1 \in (\Theta_1 \setminus \Theta_1^*)$. This means that we can find a type $\theta_2 \in \Theta_2$ such that $\phi_2(\theta_1, \theta_3^0)$ and $\phi_3(\theta_1, \theta_2)$ both fail to be interior. Recall from the previous paragraph that $\phi_1(\theta_2, \theta_3^0)$ is interior for all θ_2. Hence at the profile $(\theta_1, \theta_2, \theta_3^0)$ only agent 1 is enjoying an interior winning probability; this is impossible.

Before proceeding further, let us assign the following colors:

- red to all nodes of agent 1 with coordinates in $\Theta_2 \times \{\theta_3^0\}$. These nodes are all interior since Θ_2^∂ is empty.
- blue to all nodes of agent 2 with coordinates in $(\Theta_1 \setminus \Theta_1^*) \times \{\theta_3^0\}$. The previous paragraph implies that these nodes are all interior.
- blue to all nodes of agent 3 with coordinates in $\Theta_1^* \times \Theta_2$. These nodes are all interior by definition of Θ_1^*.

Observe that all profiles in $\Theta_1 \times \Theta_2 \times \{\theta_3^0\}$ are now either two-colored or uncolored.

If Θ_1^* is empty, then the colors assigned above already define sets R and B with the desired properties, completing the proof. Thus suppose Θ_1^* is non-empty.

Let $\theta_3 \in \Theta_3 \setminus \{\theta_3^0\}$ be arbitrary. The fact that we have already assigned blue to the nodes of agent 3 with coordinates $\Theta_1^* \times \Theta_2$ requires us to assign some colors to the nodes of agents 1 or 2 whose 3'rd coordinate is θ_3. In this step, we will not color any further nodes of agent 3. We make a case distinction.

(1) Suppose that for all θ_1 in Θ_1^* the probability $\phi_2(\theta_1, \theta_3)$ is interior. We assign red to all nodes of agent 2 with coordinates in $\Theta_1^* \times \{\theta_3\}$. This yields a coloring of the profiles in $\Theta_1 \times \Theta_2 \times \{\theta_3^0\}$ with the desired properties: The profiles in $\Theta_1^* \times \Theta_2 \times \{\theta_3\}$ are two-colored via red nodes of agent 2 and blue nodes of 3; the profiles in $(\Theta_1 \setminus \Theta_1^*) \times \Theta_2 \times \{\theta_3\}$ are uncolored.

(2) Suppose there exists $\tilde{\theta}_1 \in \Theta_1^*$ such that $\phi_2(\theta_1, \theta_3)$ is interior. Given that $\phi_3(\tilde{\theta}_1, \theta_2)$ is interior for all $\theta_2 \in \Theta_2$ (recall the definition of Θ_1^*), it must be the case that, for all $\theta_2 \in \Theta_2$, the probability $\phi_1(\theta_2, \theta_3)$ is interior.

We next claim that $\phi_2(\theta_1, \theta_3)$ is interior for all $\theta_1 \in (\Theta_1 \setminus \Theta_1^*)$. Suppose this were false for some $\theta_1 \in (\Theta_1 \setminus \Theta_1^*)$. The previous paragraph tells us that $\phi_1(\theta_2, \theta_3)$ is interior for all θ_2. Thus, if $\phi_2(\theta_1, \theta_3)$ fails to be interior, then $\phi_3(\theta_1, \theta_2)$ would have to be interior for all $\theta_2 \in \Theta_2$; this is a contradiction since $\theta_1 \in (\Theta_1 \setminus \Theta_1^*)$.

We now assign red to all nodes of agent 1 with coordinates in $\Theta_2 \times \{\theta_3\}$, and assign blue to all nodes of agent 2 with coordinates in $(\Theta_1 \setminus \Theta_1^*) \times \{\theta_3\}$. The previous two paragraphs imply that all of these nodes are interior. Moreover the profiles in $\Theta_1^* \times \Theta_2 \times \{\theta_3\}$ are two-colored via red nodes of agent 1 and blue nodes of agent 3, and the profiles in $(\Theta_1 \setminus \Theta_1^*) \times \Theta_2 \times \{\theta_3\}$ are two-colored via red nodes of agent 1 and blue nodes of agent 2.

If we apply this case distinction separately to all θ_3 in $\Theta_3 \setminus \{\theta_3^0\}$, this completes the construction of R and B in Case 2.

Case 1 and Case 2 together complete the proof of Claim 1.10. □

□

1.A.1.2 Approximate optimality of jury mechanisms

In this part of the appendix, we prove Theorem 1.4. To distinguish a random variable from its realization, we denote the former using a tilde $\sim$. Given a set N of agents, we denote the profile of their types by θ_N, and the set of these profiles by Θ_N. For example, given $i \in N$, $\omega_i \in \Omega_i$, and $\theta_{N\setminus\{i\}} \in \Theta_{N\setminus\{i\}}$, we write $\mu\left(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N\setminus\{i\}} = \theta_{N\setminus\{i\}}\right)$ to mean the probability of the event that i's value is ω_i and the types of the other agents in N are $\theta_{N\setminus\{i\}}$.

Assumption 1.2. For all $n \in \mathbb{N}$, there exists $m \in \mathbb{N}$ with the following property: Denoting $N = \{1, \dots, n\}$ and $N' = \{n+1, \dots, n+m\}$, there is a function $g\colon \Theta_{N'} \times \Theta_N \to \mathbb{R}_+$ with the following two properties:

(1) For all $i \in N$, all $\omega_i \in \Omega_i$ and $\theta_{N\setminus\{i\}} \in \Theta_{N\setminus\{i\}}$ we have

$$\begin{aligned}&\mu\big(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N\setminus\{i\}} = \theta_{N\setminus\{i\}}\big)\\ &= \sum_{\theta_{N'}\in\Theta_{N'}} \sum_{\theta_i\in\Theta_i} g(\theta_{N'}, \theta_{N\setminus\{i\}}, \theta_i)\mu\big(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N'} = \theta_{N'}\big).\end{aligned} \tag{1.A.1}$$

(2) For all $\theta_{N'} \in \Theta_{N'}$ we have

$$\sum_{\theta_N\in\Theta_N} g(\theta_{N'}, \theta_N) = 1. \tag{1.A.2}$$

Lemma 1.13. *Assumption 1.1 implies Assumption 1.2.*

Proof of Lemma 1.13. Let $m = n$. Let $N = \{1, \dots, n\}$ and $N' = \{n+1, \dots, 2n\}$, and let $\xi : N \to N'$ be a bijection. It is straightforward to verify that the function g defined as follows has the desired properties: For all $(\theta_N, \theta_{N'})$, let $g(\theta_N, \theta_{N'}) = 1$ if for all $i \in N$ the types of i and $\xi(i)$ agree; else, let $g(\theta_N, \theta_{N'}) = 0$. □

Proof of Theorem 1.4. The second part of the claim is immediate from the first. For the first part, let ϕ be an arbitrary DIC mechanism with n agents. Let $N = \{1, \dots, n\}$. For this choice of N, we invoke Lemma 1.13 to find m and g as in Assumption 1.2. Let $N' = \{n+1, \dots, n+m\}$. We define our candidate jury mechanism as follows: For all $i \in N$, let $\psi_i : \Theta_{N'} \to \mathbb{R}^n$ be defined by

$$\forall_{\theta_{N'}\in\Theta_{N'}}, \quad \psi_i(\theta_{N'}) = \sum_{\theta_N\in\Theta_N} g(\theta_{N'}, \theta_N)\phi_i^*(\theta_{N\setminus\{i\}}).$$

For all $i \in N'$, let $\psi_i = 0$. Let $\psi = (\psi_1, \dots, \psi_m)$.

Notice that ψ only depends on the reports of agents in N'. Since N' is disjoint from N, we can show that ψ is a jury mechanism in the setting with $n+m$ agents by showing that ψ maps to probability distributions over N. It is clear that ϕ is non-negative (as g and ψ^* are non-negative). To verify that ψ almost surely allocates to an agent in N, we observe that for all profiles $\theta_{N'}$ we have the following (the first equality is by definition of ψ; the second is from the fact that ϕ^* is a well-defined mechanism when the set of agents is N; the third is from (1.A.2)):

$$\sum_{i\in N}\psi_i(\theta_{N'}) = \sum_{i\in N}\sum_{\theta_N\in\Theta_N} g(\theta_{N'}, \theta_N)\phi_i^*(\theta_{N\setminus\{i\}}) = \sum_{\theta_N\in\Theta_N} g(\theta_{N'}, \theta_N) = 1,$$

as desired. We complete the proof by verifying that ϕ and ψ lead to the same expected value. We write the expected value from ϕ as follows (the first equality follows from (1.A.1); the remaining equalities obtain by rearranging):

$$
\begin{aligned}
&\sum_{i\in N}\sum_{\theta_{N\setminus\{i\}}}\sum_{\omega_i}\omega_i\mu\left(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N-i} = \theta_{N\setminus\{i\}}\right)\phi_i^*(\theta_{N\setminus\{i\}})\\
=&\sum_{i\in N}\sum_{\theta_{N\setminus\{i\}}}\sum_{\omega_i}\omega_i\sum_{\theta_{N'}}\sum_{\theta_i}g(\theta_{N'}, \theta_{N\setminus\{i\}}, \theta_i)\mu\left(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N'} = \theta_{N'}\right)\phi_i^*(\theta_{N\setminus\{i\}})\\
=&\sum_{i\in N}\sum_{\omega_i}\sum_{\theta_{N'}}\omega_i\mu\left(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N'} = \theta_{N'}\right)\sum_{\theta_{N\setminus\{i\}}}\sum_{\theta_i}g(\theta_{N'}, \theta_{N\setminus\{i\}}, \theta_i)\phi_i^*(\theta_{N\setminus\{i\}})\\
=&\sum_{i\in N}\sum_{\omega_i}\sum_{\theta_{N'}}\omega_i\mu\left(\tilde{\omega}_i = \omega_i, \tilde{\theta}_{N'} = \theta_{N'}\right)\psi_i(\theta_{N'}).
\end{aligned}
$$

This last expression is precisely the expected value from ψ. □

1.A.2 Random allocations

1.A.2.1 Proof of Proposition 1.6

Proof of Proposition 1.6. To keep calculations readable, it will be convient to adopt the following notation: When a DIC mechanism ϕ is given, we denote

$$
\begin{aligned}
&\phi_1(\theta^a) = p^{a|b}, \quad \phi_3(\theta^c) = p^{b|c}, \quad \phi_2(\theta^c) = p^{c|d}, \quad \phi_3(\theta^e) = p^{d|e},\\
&\phi_1(\theta^e) = p^{e|f}, \quad \phi_3(\theta^f) = p^{f|g}, \quad \phi_2(\theta^a) = p^{g|a}.
\end{aligned}
$$

The probabilities in the previous display do not fully describe the mechanism, but these are the only ones needed to evaluate the mechanism. For a given value of ρ, we denote the expected value from ϕ by $V_\rho(\phi)$. Direct computation shows

$$
V_\rho(\phi) = p^{a|b} + p^{b|c} + p^{c|d} + 2p^{d|e} + p^{e|f} + p^{f|g} + p^{g|a} - \rho\left(p^{b|c} + 2p^{d|e} + p^{f|g}\right). \tag{1.A.3}
$$

In particular, $V_\rho(\phi^*) = 4 - 2\rho$.

We first show that ϕ^* is uniquely optimal if $\rho \in (0, \frac{1}{2})$. The following auxiliary claim is central.

Claim 1.14. *Let ϕ be a DIC mechanism distinct from ϕ^*. We have $V_{\frac{1}{2}}(\phi) \leq V_{\frac{1}{2}}(\phi^*)$. Further, there exists $\rho_\phi \in (0, \frac{1}{2})$ such that $\rho \in (0, \rho_\phi)$ implies $V_\rho(\phi) < V_\rho(\phi^*)$.*

Proof of Claim 1.14. Inspection of Figure 1.1 shows that ϕ must satisfy the following system of inequalities:

$$
\begin{aligned}
&p^{a|b} + p^{g|a} \leq 1, \quad p^{a|b} + p^{b|c} \leq 1, \quad p^{c|d} + p^{b|c} \leq 1, \quad p^{c|d} + p^{d|e} \leq 1,\\
&p^{e|f} + p^{d|e} \leq 1, \quad p^{e|f} + p^{f|g} \leq 1, \quad p^{g|a} + p^{f|g} \leq 1.
\end{aligned} \tag{1.A.4}
$$

Turning to the first part of the claim, we have to show $V_{\frac{1}{2}}(\phi) \leq V_{\frac{1}{2}}(\phi^*)$. Direct computation shows $V_{\frac{1}{2}}(\phi^*) = 3$. Using (1.A.4), we can bound $V_{\frac{1}{2}}(\phi)$ as follows.

$$\begin{aligned}
V_{\frac{1}{2}}(\phi) &= p^{a|b} + p^{b|c} + p^{c|d} + 2p^{d|e} + p^{e|f} + p^{f|g} + p^{g|a} - \frac{1}{2}\left(p^{b|c} + 2p^{d|e} + p^{f|g}\right) \\
&= p^{a|b} + \frac{p^{b|c}}{2} + p^{c|d} + p^{d|e} + p^{e|f} + \frac{p^{f|g}}{2} + p^{g|a} \\
&= \underbrace{p^{a|b} + p^{g|a}}_{\leq 1} + \underbrace{\frac{p^{b|c} + p^{c|d}}{2}}_{\leq \frac{1}{2}} + \underbrace{\frac{p^{c|d} + p^{d|e}}{2}}_{\leq \frac{1}{2}} + \underbrace{\frac{p^{d|e} + p^{e|f}}{2}}_{\leq \frac{1}{2}} + \underbrace{\frac{p^{e|f} + p^{f|g}}{2}}_{\leq \frac{1}{2}} \\
&\leq 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} \\
&= 3.
\end{aligned}$$

Hence $V_{\frac{1}{2}}(\phi) \leq V_{\frac{1}{2}}(\phi^*)$, as promised.

Now consider the second part of the claim. We show the contrapositive: If there exists a sequence $\{\rho_k\}_{k\in\mathbb{N}}$ in $(0, \frac{1}{2})$ that converges to 0 and such that $V_{\rho_k}(\phi) \geq V_{\rho_k}(\phi^*)$ holds for all k, then $\phi = \phi^*$. Let $\{\rho_k\}_{k\in\mathbb{N}}$ be such a sequence. For all ρ_k, the system (1.A.4) implies the following upper bound on $V_{\rho_k}(\phi)$:

$$\begin{aligned}
V_{\rho_k}(\phi) &= \underbrace{p^{a|b} + p^{b|c}}_{\leq 1} + \underbrace{p^{c|d} + p^{d|e}}_{\leq 1} + \underbrace{p^{d|e} + p^{e|f}}_{\leq 1} + \underbrace{p^{f|g} + p^{g|a}}_{\leq 1} \\
&\quad - \rho_k\left(p^{b|c} + 2p^{d|e} + p^{f|g}\right) \\
&\leq 4 - \rho_k\left(p^{b|c} + 2p^{d|e} + p^{f|g}\right).
\end{aligned} \tag{1.A.5}$$

Since $V_{\rho_k}(\phi) \geq V_{\rho_k}(\phi^*) = 4 - 2\rho_k$ and $\rho_k > 0$, we find

$$p^{b|c} + 2p^{d|e} + p^{f|g} \;\leq\; 2. \tag{1.A.6}$$

Further, since $V_{\rho_k}(\phi) \geq 4 - 2\rho_k$ holds for all k, taking limits implies $V_0(\phi) \geq 4$. Together with the bound in (1.A.5) we get $V_0(\phi) = 4$; that is,

$$V_0(\phi) \;=\; p^{a|b} + p^{b|c} + p^{c|d} + p^{d|e} + p^{d|e} + p^{e|f} + p^{f|g} + p^{g|a} \;=\; 4 \tag{1.A.7}$$

Hence (1.A.4) and (1.A.7) imply

$$p^{a|b} + p^{b|c} \;=\; p^{c|d} + p^{d|e} \;=\; p^{d|e} + p^{e|f} \;=\; p^{f|g} + p^{g|a} \;=\; 1. \tag{1.A.8}$$

We now bound $V_0(\phi)$ a second time (the equality is by direct computation; the inequality follows from (1.A.4)):

$$V_0(\phi) \;=\; p^{a|b} + p^{g|a} + p^{b|c} + p^{c|d} + 2p^{d|e} + p^{e|f} + p^{f|g} \;\leq\; 3 + 2p^{d|e}. \tag{1.A.9}$$

Hence $V_0(\phi) = 4$ implies $p^{d|e} \geq \frac{1}{2}$. We next claim $p^{d|e} = \frac{1}{2}$. Towards a contradiction, suppose not, meaning $p^{d|e} > \frac{1}{2}$. Hence (1.A.8) implies $p^{c|d} = p^{e|f} < \frac{1}{2}$. Now, we also know from (1.A.6) and (1.A.7) that

$$p^{a|b} + p^{c|d} + p^{e|f} + p^{g|a} \;\geq\; 2$$

holds. However, in light of (1.A.4) we have $p^{a|b} + p^{g|a} \leq 1$, and hence the previous display requires $p^{c|d} + p^{e|f} \geq 1$. This contradicts $p^{c|d} = p^{e|f} < \frac{1}{2}$. Thus $p^{d|e} = \frac{1}{2}$.

Let us now return to the bound derived in (1.A.9). In view of $p^{d|e} = \frac{1}{2}$ and (1.A.4), we can infer from (1.A.9) that $p^{a|b} + p^{g|a} = p^{b|c} + p^{c|d} = p^{e|f} + p^{f|g} = 2p^{d|e} = 1$ holds. Together with (1.A.8), we find

$$p^{a|b} \;=\; 1 - p^{b|c} \;=\; p^{c|d} \;=\; 1 - p^{d|e} \;=\; p^{e|f} \;=\; 1 - p^{f|g} \;=\; p^{g|a}. \tag{1.A.10}$$

We already know that $p^{d|e} = \frac{1}{2}$ holds. Hence all probabilities (1.A.10) must equal $\frac{1}{2}$. This shows that ϕ agrees with ϕ^* at all profiles in $\Theta^* = \{\theta^a, \theta^b, \theta^c, \theta^d, \theta^e, \theta^f, \theta^g\}$. By inspecting $\Theta \setminus \Theta^*$, it is now easy to verify that ϕ and ϕ^* also agree on $\Theta \setminus \Theta^*$. □

We next use Claim 1.14 to show that ϕ^* is uniquely optimal if $\rho \in (0, \frac{1}{2})$. Let ϕ be an arbitrary DIC mechanisms distinct from ϕ^*. Inspection of (1.A.3) shows that the difference $V_\rho(\phi) - V_\rho(\phi^*)$ is an affine function of ρ. That is, there exist reals a_ϕ and b_ϕ such that $V_\rho(\phi) - V_\rho(\phi^*) = a_\phi + b_\phi \rho$ holds for all $\rho \in [0, \frac{1}{2}]$. Let $\rho_\phi \in (0, \frac{1}{2})$ be as in the conclusion of Claim 1.14. If $\rho \in (0, \rho_\phi)$, the choice of ρ_ϕ implies $V_\rho(\phi) < V_\rho(\phi^*)$, and so we are done. Hence in what follows we assume $\rho \in [\rho_\phi, \frac{1}{2})$. We distinguish two cases.

(1) If $b_\phi \leq 0$, then

$$V_\rho(\phi) - V_\rho(\phi^*) \;=\; a_\phi + b_\phi \rho \;\leq\; a_\phi + b_\phi \frac{\rho_\phi}{2} \;=\; V_{\frac{\rho_\phi}{2}}(\phi) - V_{\frac{\rho_\phi}{2}}(\phi^*).$$

Now $\frac{\rho_\phi}{2} \in (0, \rho_\phi)$ and the choice of ρ_ϕ imply $V_{\frac{\rho_\phi}{2}}(\phi) - V_{\frac{\rho_\phi}{2}}(\phi^*) < 0$, and we are done.

(2) If $b_\phi > 0$, then

$$V_\rho(\phi) - V_\rho(\phi^*) \;=\; a_\phi + b_\phi \rho \;<\; a_\phi + b_\phi \frac{1}{2} \;=\; V_{\frac{1}{2}}(\phi) - V_{\frac{1}{2}}(\phi^*).$$

Now Claim 1.14 implies $V_{\frac{1}{2}}(\phi) - V_{\frac{1}{2}}(\phi^*) \leq 0$, and we are done.

Hence all $\rho \in (0, \frac{1}{2})$ and all DIC mechanisms ϕ distinct from ϕ^* satisfy $V_\rho(\phi) < V_\rho(\phi^*)$.

It remains to show that ϕ^* is not uniquely optimal if $\rho \in \{0, \frac{1}{2}\}$, and that ϕ^* is not optimal if $\rho \notin [0, \frac{1}{2}]$. To that end, recall the constant mechanism and the jury mechanism described in the paragraphs after Proposition 1.6. By direct computation one can show that the constant mechanism or the jury mechanism, respectively, generate the same expected value as ϕ^* if $\rho = 0$ or $\rho = \frac{1}{2}$, respectively. Thus ϕ^* is not uniquely optimal if $\rho \in \{0, \frac{1}{2}\}$. Since ϕ^* is uniquely optimal on $(0, \frac{1}{2})$, and since the expected value is affine in ρ, we conclude that ϕ^* is not optimal if $\rho \notin [0, \frac{1}{2}]$. □

1.A.2.2 Proof of Theorem 1.5

Lemma 1.15. *If for all agents i we have $|\Theta_i| \leq 2$, then all extreme points of the set of DIC mechanisms are deterministic.*

For the proof, recall the following definitions for a given (simple undirected) graph G with node set V and edge set E. Given a node v, the set of edges which are incident to v is denoted $E(v)$. A *perfect matching* is a function $\psi: E \to \{0,1\}$ such that all $v \in V$ satisfy $\sum_{e \in E(v)} \psi(e) = 1$. The *perfect matching polytope* is the set $\{\psi: E \to [0,1]\colon \forall_{v \in V}, \sum_{e \in E(v)} \psi(e) = 1\}$.

Proof of Lemma 1.15. Let us relabel types such that we have $\Theta_i \subseteq \{0,1\}$ for all i. First, suppose we have $\Theta_i = \{0,1\}$ for all i.

For all DIC mechanisms ϕ, all agents i and all profiles θ, we may drop i's report from i's winning probability, writing $\phi_i(\theta_{-i})$ instead of $\phi_i(\theta)$. Under this convention, we claim that the set of DIC mechanisms is the perfect matching polytope of the graph G that has node set $\{0,1\}^n$ and where two nodes are adjacent if and only if they differ in exactly one coordinate. (This graph is known as the *n-hypercube*.) Indeed, each node of the graph is a type profile θ, and each edge may be identified with a pair of the form (i, θ_{-i}). The set of edges incident to θ is the set $\{(i, \theta_{-i})\}_{i=1}^n$. Hence the constraint $\sum_{e \in E(v)} \psi(e) = 1$ is exactly the constraint that the object be allocated to one of the agents.

Now, the graph G described in the previous paragraph is bi-partite (partition the type profiles (that is, the nodes of G) according to whether the profile has an odd or even number of entries equal to 0). It follows from Theorem 11.4 of Korte and Vygen (2018) that all extreme points of the perfect matching polytope are perfect matchings. All perfect matchings represent deterministic DIC mechanisms. Hence all extreme points of the set of DIC mechanisms are deterministic.

The claim for the general case, where we have $\Theta_i \subseteq \{0,1\}$ for all i, follows from the previous paragraph by viewing a DIC mechanism on Θ as a mechanism on $\{0,1\}^n$ that ignores the reports of those whose type spaces are singletons. □

Lemma 1.16. *If $|\{i \in \{1,\dots,n\}: |\Theta_i| \geq 2\}| \leq 2$, then all extreme points of the set of DIC mechanisms are deterministic.*

Proof of Lemma 1.16. We may assume $n \geq 3$, as otherwise the claim follows from Lemma 1.15. We will prove the claim for the case where $|\{i \in \{1,\dots,n\}: |\Theta_i| \geq 2\}| = 2$, the other cases being simpler. After possibly relabelling the agents, suppose we have $|\Theta_1| \geq 2$ and $|\Theta_2| \geq 2$. Let ϕ be a stochastic DIC mechanism. Notice that at all profiles θ where either agent 1 or agent 2 but not both is enjoying an interior winning probability, there must be an agent in $\{3,\dots,n\}$ who is also enjoying an interior winning probability; let i_θ denote one such agent. For a number $\varepsilon > 0$ to be chosen later, consider $f: \Theta \to \{-\varepsilon, 0, \varepsilon\}^n$ defined for all θ as follows:

(1) If $\phi_1(\theta) \in (0,1)$ and $\phi_2(\theta) \in (0,1)$, let $f_1(\theta) = \varepsilon$, let $f_2(\theta) = -\varepsilon$, and let $f_i(\theta) = 0$ for all $i \notin \{1,2\}$.

(2) If $\phi_1(\theta) \in (0,1)$ and $\phi_2(\theta) \notin (0,1)$, let $f_1(\theta) = \varepsilon$, let $f_{i_\theta}(\theta) = -\varepsilon$, and let $f_i(\theta) = 0$ for all $i \notin \{1, i_\theta\}$.

(3) If $\phi_1(\theta) \notin (0,1)$ and $\phi_2(\theta) \in (0,1)$, let $f_2(\theta) = -\varepsilon$, let $f_{i_\theta}(\theta) = \varepsilon$, and let $f_i(\theta) = 0$ for all $i \notin \{2, i_\theta\}$.

Using that, for all θ, agent i_θ has a singleton type space, it is easy to see that $\phi + f$ and $\phi - f$ are two DIC mechanisms distinct from ϕ whenever ε is sufficiently small. Thus ϕ is not an extreme point. □

Proof of Theorem 1.5. Lemma 1.3, Lemma 1.15 and Lemma 1.16 together imply that all extreme points are deterministic if one of the conditions (1) to (3) holds. Now let conditions (1) to (3) all fail. We know from Section 1.5.2 that a stochastic extreme point exists in the hypothetical situation where $n = 4$ and the set of type profiles is $\hat{\Theta} = \{\ell, r\} \times \{u, d\} \times \{f, c, b\} \times \{0\}$. Since (1) to (3) all fail, we can relabel the agents and types such that agents 1 to 4 have these sets as subsets of their respective sets of types. Let ϕ^* denote the stochastic extreme point Section 1.5.2. Using ϕ^*, it is straightforward to define a stochastic extreme point for the actual set of type profiles with n agents. To see this in detail, let us agree to the following notation: when $i \in \{1,2,3\}$, then $\hat{\Theta}_{-i}$ means the sets of type profiles of agents $\{1,2,3,4\} \setminus \{i\}$ that belong to $\hat{\Theta}$. Now consider $\psi^* : \Theta \to \mathbb{R}^n$ defined as follows: For all $i \in \{1,\dots,n\} \setminus \{1,2,3,4\}$, let $\psi_i^* = 0$; for all $i \in \{1,2,3\}$ and all $\theta \in \Theta$, let $\psi_i^*(\theta) = \phi_i^*(\theta_1, \theta_2, \theta_3, \theta_4)$ if $(\theta_j)_{j \in \{1,2,3,4\} \setminus \{i\}} \in \hat{\Theta}_{-i}$, and let $\psi_i^*(\theta) = 0$ if $(\theta_j)_{j \in \{1,2,3,4\} \setminus \{i\}} \notin \hat{\Theta}_{-i}$; let $\psi_4^* = 1 - \sum_{i=1}^{3} \psi_i^*$. A moment's thought reveals that ψ^* is a well-defined DIC mechanism. To see that it is a stochastic extreme point, consider an arbitrary DIC mechanism ψ that appears in a convex combination that equals ψ^*. We know from Section 1.5.2 that ψ must agree with ψ^* whenever the types of agents 1 to 4 are in $\hat{\Theta}$. From here it is easy to see that ψ must agree with ψ^* at all other profiles, too. □

1.A.3 Anonymous juries

1.A.3.1 Proof of Theorem 1.7

Proof of Theorem 1.7. Let ϕ be DIC and anonymous.

The following notation is useful. Let T denote the common type space. Let T^{n-1} with generic element θ^{n-1} denote the $(n-1)$-fold Cartesian product of T. We will frequently consider profiles obtained from a profile θ^{n-1} in T^{n-1} by replacing one entry of θ^{n-1}. For instance, we write (t, θ_{-j}^{n-1}) to denote the profile obtained by replacing the j'th entry of θ^{n-1} by t.

By DIC, for all i, we may drop i's type from i's winning probability. Thus we write $\phi_i(\theta^{n-1})$ for i's winning probability when the types of the others are $\theta^{n-1} \in T^{n-1}$. Anonymity implies that $\phi_i(\theta^{n-1})$ is invariant to permutations of θ^{n-1}.

We use the following auxiliary claim.

Claim 1.17. *Let $i \in \{1, \dots, n\}$, $t \in T$, $t' \in T$, and $\theta^{n-1} \in T^{n-1}$. Then*

$$\sum_{j=1}^{n-1} \Big(\phi_i(t, \theta^{n-1}_{-j}) - \phi_i(t', \theta^{n-1}_{-j}) \Big) = 0. \tag{1.A.11}$$

Proof of Claim 1.17. Let us arbitrarily label θ^{n-1} as $(\theta_j)_{j \in N \setminus \{i\}}$. Let us also fix an arbitrary type $\theta_i \in T$.

In an intermediate step, let j be distinct from i. For clarity, we spell out winning probabilities as follows: $\phi_i(r_i = t, r_j = t', r_{-ij} = \theta_{-ij})$ means i's winning probability when i reports t, j reports t', and all remaining agents report θ_{-ij}. A permutation of i's and j's reports does not change the winning probabilities of the agents other than i and j. Since the object is allocated with probability one, we have

$$\begin{aligned} &\phi_i(r_i = t, r_j = t', r_{-ij} = \theta_{-ij}) + \phi_j(r_i = t, r_j = t', r_{-ij} = \theta_{-ij}) \\ =&\phi_i(r_i = t', r_j = t, r_{-ij} = \theta_{-ij}) + \phi_j(r_i = t', r_j = t, r_{-ij} = \theta_{-ij}). \end{aligned}$$

By rearranging the previous display, and by DIC, we obtain

$$\begin{aligned} &\phi_i(r_i = t, r_j = t', r_{-ij} = \theta_{-ij}) - \phi_i(r_i = t', r_j = t, r_{-ij} = \theta_{-ij}) \\ =&\phi_j(r_i = t', r_j = \theta_j, r_{-ij} = \theta_{-ij}) - \phi_j(r_i = t, r_j = \theta_j, r_{-ij} = \theta_{-ij}). \end{aligned} \tag{1.A.12}$$

Now consider summing (1.A.12) over all $j \in \{1, \dots, n\} \setminus \{i\}$. This summation yields

$$\sum_{j:\, j \neq i} \big(\phi_i(r_i = t, r_j = t', r_{-ij} = \theta_{-ij}) - \phi_i(r_i = t', r_j = t, r_{-ij} = \theta_{-ij}) \big) \tag{1.A.13}$$

$$= \sum_{j:\, j \neq i} \big(\phi_j(r_i = t', r_j = \theta_j, r_{-ij} = \theta_{-ij}) - \phi_j(r_i = t, r_j = \theta_j, r_{-ij} = \theta_{-ij}) \big). \tag{1.A.14}$$

In (1.A.14), the profiles considered are all of the form $(r_i = t', r_{-i} = \theta_{-i})$ and $(r_i = t, r_{-i} = \theta_{-i})$, respectively. Note that by DIC we have $\phi_i(r_i = t', r_{-i} = \theta_{-i}) - \phi_i(r_i = t, r_{-i} = \theta_{-i}) = 0$. Hence (1.A.14) equals

$$\sum_{j=1}^{n} \big(\phi_j(r_i = t', r_{-i} = \theta_{-i}) - \phi_j(r_i = t, r_{-i} = \theta_{-i}) \big).$$

Since the object is always allocated, the term in the previous display equals 0. Hence the sum in (1.A.13) equals

$$\sum_{j:\, j\neq i} \Big(\phi_i(r_i = \theta_i, r_j = t', r_{-ij} = \theta_{-ij}) - \phi_i(r_i = \theta_i, r_j = t, r_{-ij} = \theta_{-ij})\Big) = 0.$$

We now revert to our usual notation. By DIC, we may drop i's report from ϕ_i. Since ϕ_i is permutation-invariant with respect to $N \setminus \{i\}$, we may also write

$$\phi_i(r_i = \theta_i, r_j = t', r_{-ij} = \theta_{-ij}) = \phi_i(t', \theta_{-j}^{n-1}) \qquad \text{and}$$
$$\phi_i(r_i = \theta_i, r_j = t, r_{-ij} = \theta_{-ij}) = \phi_i(t, \theta_{-j}^{n-1}).$$

Thus we obtain the desired equality $\sum_{j=1}^{n-1} \Big(\phi_i(t', \theta_{-j}^{n-1}) - \phi_i(t, \theta_{-j}^{n-1})\Big) = 0.$ □

In what follows, let i be an arbitrary agent. We show i's winning probability is constant in the reports of others. To that end, let us fix an arbitrary type $t^* \in T$. For all $k \in \{0, \ldots, n-1\}$, let T_k^{n-1} denote the subset of profiles in T^{n-1} where exactly k-many entries are distinct from t^*. Let p_i denote i's winning probability when all other agents report t^*. We will show via induction over k that i's winning probability is equal to p_i whenever the others report a profile in T_k^{n-1}. This completes the proof since $T^{n-1} = \cup_{k=0}^{n-1} T_k^{n-1}$ holds.

Base case $k = 0$. Immediate from the definitions of p_i and T_0^{n-1}.

Induction step. Let $k \geq 1$. Let all $\hat{\theta}^{n-1} \in \cup_{\ell=0}^{k-1} T_\ell^{n-1}$ satisfy $\phi_i(\hat{\theta}^{n-1}) = p_i$. Letting $\theta^{n-1} \in T_k^{n-1}$ be arbitrary, we show $\phi_i(\theta^{n-1}) = p_i$.

By anonymity, we may assume that exactly the first k entries of θ^{n-1} are distinct from t^*. That is, there exist types $t_1, \ldots, t_k$ all distinct from t^* such that $\theta^{n-1} = (t_1, \ldots, t_k, t^*, \ldots, t^*)$.

Let $\tilde{\theta}^{n-1} = (t_1, \ldots, t_{k-1}, t^*, \ldots, t^*)$. This profile is obtained from θ^{n-1} by replacing t_k by t^*. We now invoke Claim 1.17 to infer

$$\sum_{j=1}^{n-1} \phi_i(t_k, \tilde{\theta}_{-j}^{n-1}) = \sum_{j=1}^{n-1} \phi_i(t^*, \tilde{\theta}_{-j}^{n-1}). \tag{1.A.15}$$

Consider the profiles appearing in the sum on the left of (1.A.15) as j varies from 1 to $n-1$.

(1) Let $j \leq k-1$. Since exactly the first $k-1$ entries of $\tilde{\theta}$ are distinct from t^*, it follows that $(t_k, \tilde{\theta}_{-j}^{n-1})$ is another profile where exactly $k-1$ entries differ from t^*. Hence the induction hypothesis implies $\phi_i(t_k, \tilde{\theta}_{-j}^{n-1}) = p_i$.

(2) Let $j > k-1$. In the profile $(t_k, \tilde{\theta}_{-j}^{n-1})$, the first $k-1$ entries are $t_1, \ldots, t_{k-1}$, the j'th entry is t_k, and all remaining entries are t^*. Hence $(t_k, \tilde{\theta}_{-j}^{n-1})$ is a permutation of θ^{n-1}. Anonymity implies $\phi_i(t_k, \tilde{\theta}_{-j}^{n-1}) = \phi_i(\theta^{n-1})$.

Hence the sum on the left of (1.A.15) equals $\sum_{j=1}^{n-1} \phi_i(t, \tilde{\theta}_{-j}^{n-1}) = (k-1)p_i + (n-k)\phi_i(\theta^{n-1})$

Now consider the sum on the right of (1.A.15). For all j, a moment's thought reveals that the profile $(t^*, \tilde{\theta}_{-j}^{n-1})$ contains at most $(k-1)$-many entries different from t^*. By the induction hypothesis, therefore, the sum on the right of (1.A.15) equals $(n-1)p_i$.

The previous two paragraphs and (1.A.15) imply $(k-1)p_i + (n-k)\phi_i(\theta^{n-1}) = (n-1)p_i$. Equivalently, $(n-k)(\phi_i(\theta^{n-1}) - p_i) = 0$. Since $k \leq n-1$, we find $\phi_i(\theta^{n-1}) = p_i$, as promised. $\square$

1.A.3.2 Proof of Theorem 1.8

Proof of Theorem 1.8. We omit the straightforward verification that a jury mechanism with an anonymous jury is partially anonymous.

For the converse, let ϕ be deterministic, partially anonymous, and DIC. Let N denote the set of agents, and let T denote the common type space. For this proof, we write $\phi(\theta)$ to mean the agent who wins at profile θ; this makes sense since ϕ is deterministic.

Let I_i denote the set of agents that influence agent i's winning probability. For all $j \in N$, let $A_j = \{i \in N : j \in I_i\}$ be the set of agents that are influenced by j. Let $I = \{i \in N : A_i \neq \emptyset\}$. We may assume that ϕ is non-constant, meaning $I \neq \emptyset$, as otherwise the proof is trivial.

Given two agents i and j, let $D_{i-j} = A_i \setminus A_j$, and $D_{j-i} = A_j \setminus A_i$, and $C_{ij} = A_j \cap A_i$, and $N_{ij} = N \setminus (A_i \cup A_j)$. Note that, by DIC, the set C_{ij} contains neither i nor j. Hence partial anonymity implies that for all $k \in C_{ij}$ the winning probability of k is invariant with respect to permutations of i and j.

When i, j, and k are given, we write $(t, t', t'', \theta_{-ijk})$ to mean the profile where i, j, and k, respectively, report t, t', and t'', respectively, and all others report θ_{-ijk}.

Claim 1.18. *Let i and j be distinct. Let $\theta_{-ij} \in \Theta_{-ij}$. If there exists $\theta_i, \theta_j \in T$ such that $\phi(\theta_i, \theta_j, \theta_{-ij}) \in D_{i-j}$, then all $\theta_i', \theta_j' \in T$ satisfy $\phi(\theta_i', \theta_j', \theta_{-ij}) \in D_{i-j}$.*

Proof of Claim 1.18. We drop the fixed type profile θ_{-ij} of the others from the notation. To show $\phi(\theta_i', \theta_j') \in D_{i-j}$, it suffices to show $\phi(\theta_i', \theta_j) \in D_{i-j}$ since if the latter is true then definition of D_{i-j} implies $\phi(\theta_i', \theta_j') = \phi(\theta_i', \theta_j)$.

We first claim $\phi(\theta_j, \theta_i) \in D_{i-j}$. If $\phi(\theta_j, \theta_i) \in N_{ij}$, then $\phi(\theta_j, \theta_i) = \phi(\theta_i, \theta_j)$, and we have a contradiction to $\phi(\theta_i, \theta_j) \in D_{i-j}$. If $\phi(\theta_j, \theta_i) \in C_{ij}$, then partial anonymity implies $\phi(\theta_i, \theta_j) \in C_{ij}$, and we have another contradiction to $\phi(\theta_i, \theta_j) \in D_{i-j}$. If $\phi(\theta_j, \theta_i) \in D_{j-i}$, then $\phi(\theta_j, \theta_i) = \phi(\theta_i, \theta_i) \in D_{j-i}$. However, from $\phi(\theta_i, \theta_j) \in D_{i-j}$ we know $\phi(\theta_i, \theta_j) = \phi(\theta_i, \theta_i) \in D_{i-j}$; contradiction. Thus $\phi(\theta_j, \theta_i) \in D_{i-j}$.

We next claim $\phi(\theta_i', \theta_j) \in (D_{i-j} \cup C_{ij})$. Towards a contradiction, suppose not. Then $\phi(\theta_i', \theta_j) \in (D_{j-i} \cup N_{ij})$, and hence $\phi(\theta_i', \theta_j) = \phi(\theta_i, \theta_j) \notin D_{i-j}$. This contradicts the assumption $\phi(\theta_i, \theta_j) \in D_{i-j}$.

In view of the previous paragraph, we can complete the proof by showing $\phi(\theta_i', \theta_j) \notin C_{ij}$. Towards a contradiction, let $\phi(\theta_i', \theta_j) \in C_{ij}$. Partial anonymity implies $\phi(\theta_j, \theta_i') \in C_{ij}$. We have shown earlier that $\phi(\theta_j, \theta_i) \in D_{i-j}$ holds. Hence $\phi(\theta_j, \theta_i') \in D_{i-j}$, and this contradicts $\phi(\theta_j, \theta_i') \in C_{ij}$. Thus $\phi(\theta_i', \theta_j) \notin C_{ij}$ and the proof is complete. □

Claim 1.19. *Let i,j,k be distinct. Let $\theta_k \in T$ and $\theta_{-ijk} \in \Theta_{-ijk}$ be such that all $\theta_i', \theta_j' \in T$ satisfy $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in (C_{ij} \cup N_{ij})$. Then, all $\theta_i', \theta_j', \theta_k' \in T$ satisfy $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in (C_{ij} \cup N_{ij})$.*

Proof of Claim 1.19. Towards a contradiction, suppose $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in (D_{i-j} \cup D_{j-i})$. Suppose $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in D_{i-j}$, the other case being similar. The inclusions $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in (C_{ij} \cup N_{ij})$ and $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in D_{i-j}$ together imply $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in A_k$. Hence $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in D_{k-j}$. We now invoke Claim 1.18 to infer $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in D_{k-j}$. Since we also have $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in (C_{ij} \cup N_{ij})$, we infer $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in N_{ij}$. In particular, we have $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \notin A_i$. Hence $\phi(\theta_i', \theta_j', \theta_k, \theta_{-ijk}) \in D_{k-i}$. We now invoke Claim 1.18 to infer $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in D_{k-i}$. In particular, we have $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \notin A_i$. This contradicts the assumption $\phi(\theta_i', \theta_j', \theta_k', \theta_{-ijk}) \in D_{i-j}$. □

Claim 1.20. *If $C_{ij} \neq \emptyset$, then $D_{i-j} \cup D_{j-i} = \emptyset$.*

Proof of Claim 1.20. Let $k \in C_{ij}$. We may find a profile θ such that $\phi(\theta) = k$ as else k's winning probability is constantly 0 (which would contradict $k \in C_{ij}$). Denoting by θ_{-ij} the types of agents other than i and j at θ, we appeal to Claim 1.18 to infer that all $\theta_i', \theta_j' \in T$ satisfy $\phi(\theta_i', \theta_j', \theta_{-ij}) \in (C_{ij} \cup N_{ij})$. Repeatedly applying Claim 1.19 implies that all profiles θ' satisfy $\phi(\theta') \in (C_{ij} \cup N_{ij})$. It follows that all agents in $D_{i-j} \cup D_{j-i}$ enjoy a winning probability that is constantly equal to 0. Recalling the definitions $D_{i-j} = A_i \setminus A_j$, and $D_{j-i} = A_j \setminus A_i$, it follows that $D_{i-j} \cup D_{j-i}$ is empty. □

Recall the definition $I = \{i \in N : A_i \neq \emptyset\}$. Consider the binary relation $\sim$ on I defined as follows: Given i and j in I, we let $i \sim j$ if and only if $C_{ij} \neq \emptyset$.

Claim 1.21. *The relation $\sim$ is an equivalence relation. For all $i,j \in I$, if $i \sim j$, then $i \notin A_j$ and $A_i = A_j$.*

Proof of Claim 1.21. It is clear that $\sim$ is symmetric. As for reflexivity, note that $i \in I$ implies $A_i = C_{ii} \neq \emptyset$. Turning to transitivity, suppose $i \sim j$ and $j \sim k$. Hence $C_{ij} \neq \emptyset$ and $C_{jk} \neq \emptyset$. Let $\ell \in C_{jk}$. Claim 1.20 and $C_{ij} \neq \emptyset$ together imply $D_{j-i} = \emptyset$. Hence $\ell \in C_{jk}$ implies $\ell \in C_{ij}$. Hence $\ell \in C_{jk} \cap C_{ij}$, implying $\ell \in C_{ik}$. Hence $i \sim k$.

As for the second part of the claim, let $i \sim j$. Thus $C_{ij} \neq \emptyset$. Claim 1.20 implies $D_{j-i} = D_{i-j} = \emptyset$. This immediately implies $A_i = A_j$. Together with DIC, we also infer $i \notin A_j$. □

Claim 1.21 implies that we may partition I into finitely-many non-empty $\sim$-equivalence classes. (Recall that I is non-empty.) We now claim that there is exactly one $\sim$-equivalence class. Towards a contradiction, suppose not. In view of Claim 1.21, this means that there are distinct i and j such that $A_i \cap A_j = \emptyset$ and $A_i \neq \emptyset \neq A_j$. Let J_i and J_j, respectively, denote the equivalence classes containing i and j, respectively. Let $k \in A_i$ and $\ell \in A_j$. Claim 1.21 implies $k \notin J_i$ and $\ell \notin J_j$ and $k \neq \ell$. Since $k \in A_i$ and ϕ is deterministic, there is a type profile θ such that $\phi(\theta) = k$; there must be another type profile θ' such that $\phi(\theta') = \ell$. However, the definition of equivalence classes implies that k's winning probability depends only on the types of agents in J_i, and that ℓ's winning probability depends only on the types of agents in J_j. Hence there is a type profile where both k and ℓ are winning with probability 1 (such a type profile is obtained by changing at the profile θ the types of agents in J_j to their respective types at θ', and keeping all other types fixed). Contradiction.

Now, Claim 1.21 implies that the members of the unique $\sim$-equivalence class do not influence one another, and that they influence the same set of others. By partial anonymity, it follows ϕ that is a deterministic jury mechanism with an anonymous jury. □

1.A.3.3 Proof of Proposition 1.9

We first give an informal sketch of the proof. The idea is to "symmetrize" the stochastic extreme point ϕ^* from Section 1.5.2.

In Section 1.5.2, there are four agents, the set of type profiles of agents 1 to 3 is a $2 \times 2 \times 3$ set $\hat{\Theta}$, and agent 4 has a singleton type space. Let us view allocating to agent 4 as disposing the object. Let us relabel the types of agents 1 to 3 so that they are all distinct. Across agents 1 to 3 we thus have a set T of seven distinct types. The 3-fold Cartesian product T^3 of T with itself contains six permutations of $\hat{\Theta}$ (one for each permutation of $\{1,2,3\}$). In Figure 1.A.1, the common type space is labelled $T = \{1,\ldots,7\}$, and the six permutations of $\hat{\Theta}$ are depicted via six symbols (square, circle, etc.).

We can associate to each permutation of $\hat{\Theta}$ a permutation of the mechanism ϕ^*. The idea is to extend these permutations to a DIC mechanism with disposal on all of T^3. The difficulty is to verify that the resulting mechanism is well-defined. To see the issue, reconsider Figure 1.A.1. For each of the six subsets, imagine rays emanating from the subset and travelling parallel to the axes. Along the ray, exactly one agent's type changes. Hence DIC requires that this agent's winning probability remain constant along the ray. The rays emanating from distinct subset intersect, and we have verify that the sum of the associated winning probabilities does not exceed 1. We use two observations. The first is that, at most two such rays intersect simultaneously; this is a consequence of the fact that the types in $\hat{\Theta}$ are distinct across agents. The second is that the winning probabilities associated with the rays are at most $\frac{1}{2}$; this is a consequence of the construction of ϕ^* in Section 1.5.2.

Proof of Proposition 1.9. We first prove part (2) of the claim, assuming for a moment that part (1) is true. Let $\psi^*\colon T^3 \to [0,1]^3$ be a mechanism with disposal for three agents that meets the conclusion of part (1). We view ψ^* as a mechanism (without disposal) with four agents that ignores the report of agent 4, and where agent 4's winning probability equals the probability that ψ^* does not allocate the object to the first three agents. Using the assumed properties of ψ^*, we obtain a mechanism without disposal that is DIC, partially anonymous, and an extreme point of the set of DIC mechanisms without disposal for four agents.

It remains to prove part (1) of the claim. That is, we show that if $n = 3$, then there is an anonymous DIC mechanism with disposal that is an extreme point of the set of all DIC mechanisms with disposal.

Let us relabel the common type space as $T = \{1,2,3,4,5,6,7\}$. Let $T^3 = \times_{i=1}^3 T$ denote the 3-fold Cartesian product of T. Let $T_1 = \{1,2\}$, $T_2 = \{3,4\}$ and $T_3 = \{5,6,7\}$ and $\hat{\Theta} = T_1 \times T_2 \times T_3$. In Section 1.5.2, we constructed a stochastic DIC mechanism ϕ^* without disposal in a setting with 4 agents, where the types of agents 1, 2, and 3, respectively, are $\{\ell, r\}$, $\{u, d\}$, $\{f, c, b\}$, respectively, and where agent 4's type is degenerate. By relabeling types, we view ϕ^* as a mechanism with disposal with 3 agents on the set of type profiles $\hat{\Theta}$, and where allocating to agent 4 is identified with disposing the object. The arguments from Section 1.5.2 show that, if $n = 3$ and the set of type profiles is $\hat{\Theta}$, then ϕ^* is an extreme point of the set of DIC mechanisms with disposal.

For later reference, we note that, at all type profiles $\theta \in \hat{\Theta}$ and all $i \in \{1,2,3\}$, agent i's winning probability at θ under ϕ^* is either 0 or 1/2.

Our candidate mechanism will be denoted ψ^*. Let Ξ denote the set of permutations of $\{1,2,3\}$. Let $\Theta^* = \{\xi(\theta)\colon \theta \in \hat{\Theta}, \xi \in \Xi\}$ denote the set of type profiles obtained by permuting a type profile in $\hat{\Theta}$; see Figure 1.A.1. Fixing an arbitrary type profile in $\hat{\Theta}$, the types of the agents at this type profile are all distinct. Consequently, for all θ^* in Θ^* there is a unique profile θ in $\hat{\Theta}$ and ξ in Ξ such that $\theta^* = \xi(\theta)$.

For later reference, we also note that at an arbitrary type profile in Θ^*, the types of distinct agents must belong to distinct elements of the partition $\{T_1, T_2, T_3\}$.

We define ψ^* as follows: For all θ^* in Θ^*, we find the unique $(\theta, \xi) \in T \times \Xi$ such that $\theta^* = \xi(\theta)$, and then let

$$\left(\psi_i^*(\theta^*)\right)_{i=1}^n = \left(\phi_{\xi(i)}^*(\xi(\theta))\right)_{i=1}^n. \tag{1.A.16}$$

For the remaining profiles, we proceed as follows: For all agents i and profiles θ, if θ differs from at least one profile θ^* in Θ^* in agent i's type and no other agent's

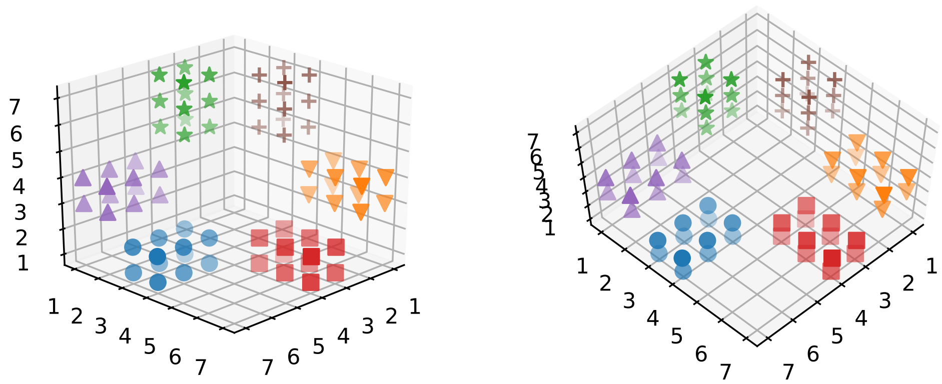

Figure 1.A.1. The set Θ^* viewed from two different angles. Each agent is associated with a distinct axis. Each symbol (square, circle, upward-pointing triangle, etc.) identifies a particular permutation of $\{1, 2, 3\}$. For instance, the upward-pointing triangles are obtained from the downward-pointing triangles by permuting the two agents on the horizontal axes.

type, then i's winning probability at θ equals i's winning probability at θ^* (which makes sense since the latter probability has already been defined in (1.A.16)); else, if no such profile θ^* in Θ^* exists, then agent i's winning probability is set equal to 0.

To complete the argument, we have to show that ψ^* is well-defined, DIC, stochastic, anonymous, and an extreme point of the set of DIC mechanisms with disposal. Assuming for a moment that the mechanism is well-defined, it is clear that the mechanism is stochastic, and one can easily verify from the definition that it is DIC and anonymous. To show that it is an extreme point of the set of DIC mechanisms, we proceed via the arguments from Section 1.5.2. Indeed, we know from Section 1.5.2 that all DIC mechanisms ψ with disposal that appear in a candidate convex combination must agree with ψ^* on $\hat{\Theta}$, and hence on Θ^*. It is then straightforward to verify that such a mechanism ψ must also agree with ψ^* on $\Theta \setminus \Theta^*$.

It remains to show that ψ^* is well-defined. Towards a contradiction, suppose there is a profile $\theta = (\theta_1, \theta_2, \theta_3)$ in Θ such that $\sum_{i=1}^{3} \psi_i^*(\theta) > 1$. By construction, all $i \in \{1, 2, 3\}$ satisfy $\psi_i^* \in \{0, \frac{1}{2}\}$. Hence all three agents enjoy non-zero winning probabilities at θ. By definition of ψ^*, we can infer the following: Since agent 1's winning probability at θ is non-zero, there exists t_1 such that $(t_1, \theta_2, \theta_3) \in \Theta^*$. Similarly, there are t_2 and t_3 such that $(\theta_1, t_2, \theta_3) \in \Theta^*$ and $(\theta_1, \theta_2, t_3) \in \Theta^*$. Recall that $\{T_1, T_2, T_3\}$ is a partition of T. Hence, for all agents i, there is a unique integer $\xi(i)$ in $\{1, 2, 3\}$ such that $\theta_i \in T_{\xi(i)}$. We now recall that if a profile is in Θ^*, then the types of distinct agents belong to distinct elements of the partition $\{T_1, T_2, T_3\}$. Hence we infer from $(t_1, \theta_2, \theta_3) \in \Theta^*$ that $\xi(2) \neq \xi(3)$ holds. Similarly, from $(\theta_1, t_2, \theta_3) \in \Theta^*$ and $(\theta_1, \theta_2, t_3) \in \Theta^*$ we infer $\xi(1) \neq \xi(2)$ and $\xi(1) \neq \xi(3)$. Taken together, we infer $\theta \in \Theta^*$. Hence the vector of winning probabilities at θ is a permutation of the vector of winning probabilities at a profile θ' in $\hat{\Theta}$. At the profile θ', the winning probabilities under ψ^* agree with ϕ^*. Thus there is a profile where the winning probabilities under ϕ^* sum to a number strictly greater than 1. This contradicts the fact that ϕ^* is a well-defined mechanism on $\hat{\Theta}$. □

Appendix 1.B Supplementary material: Disposal

In this part of the appendix, we relax the requirement that the object always be allocated. An interpretation is that the mechanism designer can dispose or privately consume the object. Accordingly, we refer to such mechanisms as mechanisms with disposal. We discuss how this affects our results from the main text (Appendix 1.B.1). Further, we show how the existence of stochastic extreme points of the set of DIC mechanisms with disposal can be related to a certain graph (Appendix 1.B.2).

Beginning with the definitions, a *mechanism with disposal* is a function $\phi : \Theta \to [0,1]^n$ satisfying

$$\forall_{\theta \in \Theta}, \quad \sum_{i=1}^{n} \phi_i(\theta) \;\le\; 1.$$

A mechanism from the main text will be referred to as a mechanism with no disposal. If there is no risk of confusion, we will drop the qualifiers "with disposal" or "with no disposal".

A mechanism with disposal is DIC if and only if for arbitrary i the winning probability ϕ_i is constant in i's report. We will sometimes drop i's report θ_i from $\phi_i(\theta_i, \theta_{-i})$.

A jury mechanism with disposal is defined as in the basic model: For all i, if agent i influences the allocation, then i never wins the object.

We normalize the value from not allocating the object to 0.

A mechanism with n agents and disposal can be viewed as a mechanism with no disposal and with $n+1$ agents where agent $n+1$ has a singleton type space; the value from allocating to $n+1$ is always 0. Likewise, if there are other agents with singleton type spaces, we can always renormalize values and view allocating to one of these agents as disposing the object. In what follows, whenever considering mechanisms with disposal, let us thus simplify by assuming that no agent has a singleton type space; that is, for all agents i we have $|\Theta_i| \ge 2$.

1.B.1 Results from the main text

Here we discuss how our results change when the mechanism can dispose the object.

To begin with, we have the following analogue of Theorem 1.5.

Theorem 1.22. *Fix n and $\Theta_1, \dots, \Theta_n$. For all agents i, let $|\Theta_i| \ge 2$. All extreme points of the set of DIC mechanisms with disposal are deterministic if and only if at least one of the following is true:*

(1) We have $n \le 2$.

(2) For all agents i we have $|\Theta_i| = 2$.

Proof of Theorem 1.22. As discussed above, a DIC mechanism with n agents and disposal is a DIC mechanism with $n+1$ agents and no disposal. The claim follows from Theorem 1.5. □

Further below, we provide an alternative proof of Theorem 1.22 that does not invoke Theorem 1.5 but relies on graph-theoretic results. We emphasize that Theorem 1.22 does not imply Theorem 1.5. Namely, we cannot conclude from Theorem 1.22 that if $n = 3$ all extreme points of the set of DIC mechanisms with no disposal are deterministic.

It follows from Theorem 1.22 that Theorem 1.1 (jury mechanisms with 3 agents) carries over to mechanisms with disposal in the sense that all mechanisms with disposal and 2 agents are convex combinations of deterministic jury mechanisms with disposal. Note that, according to Theorem 1.22, this result does not extend to $n = 3$. With $n = 2$, a jury mechanism with disposal admits a single juror whose report determines whether or not the object is disposed or allocated to the other agent.

Proposition 1.6 (on the suboptimality of deterministic DIC mechanisms) analogizes straightforwardly to mechanisms with disposal. Indeed, note that in our proof of Proposition 1.6 agent 4 was simply a dummy agent with value normalized to 0.

Theorem 1.4 (approximate optimality of jury mechanisms under Assumption 1.1 and large n) extends to mechanisms with disposal in a straightforward way, with no changes to the proof.

We already showed via Proposition 1.9 that Theorem 1.7 does not extend to mechanisms with disposal. In fact, the non-constant mechanism constructed in the proof of Proposition 1.9 actually satisfies an even stronger notion of anonymity. Namely, whenever one permutes the type profiles, the vector of winning probabilities is permuted in the same manner.

We next turn to partial anonymity for mechanisms with disposal. In particular, we show that Theorem 1.8 extends under a slight strengthening of partial anonymity. Given a mechanism ϕ, let $\phi_0 = 1 - \sum_{i=1}^{n} \phi_i$ denote the probability that the object is not allocated.

Definition 1.4. Let ϕ be a mechanism with disposal. Let $N = \{1, \ldots, n\}$ and $N_0 = N \cup \{0\}$.

(1) Given distinct $i \in N$ and $k \in N_0$, agent i *influences* k if ϕ_k is non-constant in i's report.
(2) The mechanism is *partially $*$-anonymous* if for all $i \in N$, $j \in N$, and $k \in N_0$ that are all distinct and are such that i and j influence k, agents i and j are exchangeable for k.

In words, partial anonymity is strengthened by demanding that the disposal probability ϕ_0 is permutation-invariant with respect to those agents who influence ϕ_0.

It follows from Theorem 1.8 that a deterministic partially $*$-anonymous DIC mechanism with disposal is a deterministic jury mechanism with an anonymous jury. To see this, let us view disposing the object as allocating to agent 0. Now, agent 0 does not have the same type space as the other agents. Since this was a maintained assumption of Section 1.6, we cannot yet appeal to Theorem 1.8. But, we can simply view the mechanism as a mechanism where agent 0's type space is same as the type spaces of the others, and where agent 0's report is always ignored. By now appealing to Theorem 1.8, the claim follows.

1.B.2 Stochastic extreme points and perfect graphs

In this section, we relate the existence of stochastic extreme points with disposal to a graph-theoretic property called perfection.

1.B.2.1 Preliminaries

We first recall several definitions for a simple undirected graph G with nodes V and edges E.

An *induced cycle of length* k is a subset $\{v_1, \dots, v_k\}$ of V such that, denoting $v_{k+1} = v_1$, two nodes v_ℓ and $v_{\ell'}$ in the subset are adjacent if and only if $|\ell - \ell'| = 1$.

The *line graph* of G is the graph that has as node set the edge set of G; two nodes of the line graph are adjacent if and only if the two associated edges of G share a node in G.

A *clique* of G is a set of nodes such that every pair in the set are adjacent. A clique is *maximal* if it is not a strict subset of another clique. A *stable set* of G is a subset of nodes of which no two are adjacent. The *incidence vector* of a subset of nodes $\hat{V}$ is the function $x\colon V \to \{0,1\}$ that equals one on $\hat{V}$ and equals zero otherwise. Let $S(G)$ denote the set of incidence vectors belonging to some stable set of G.

The upcoming result uses another property of graphs called *perfection*. For our purposes, it will be enough to know the following facts, all of which may be found in Korte and Vygen (2018).

Lemma 1.23. *All bi-partite graphs and line graphs of bi-partite graphs are perfect. If a graph admits an induced cycle of odd length greater than five, then it is not perfect.*

Our interest in perfect graphs is due to the following theorem of Chvátal (1975, Theorem 3.1); one may also find it in Korte and Vygen (2018, Theorem 16.21).

Theorem 1.24. *A graph G with node set V and edge set E is perfect if and only if the convex hull* $\operatorname{co} S(G)$ *is equal to the set*

$$\left\{x\colon V \to [0,1] : \textit{all maximal cliques } X \textit{ of } G \textit{ satisfy} \sum_{v \in X} x(v) \leq 1\right\}. \tag{1.B.1}$$

The set $\operatorname{co} S(G)$ is the *stable set polytope* of G. The set in (1.B.1) is the *clique-constrained stable set polytope* of G.

1.B.2.2 The feasibility graph

We next define a graph G such that the set of deterministic DIC mechanisms with disposal corresponds to $S(G)$, and such that the set of all DIC mechanisms with disposal coincides with the clique-constrained stable set polytope of G. In view of Theorem 1.24, the question of whether all extreme points are deterministic thus reduces to checking whether G is a perfect graph.

Consider the following graph G with node set V and edge set E. Let

$$V = \cup_{i=1}^{n} \left(\{i\} \times \Theta_{-i}\right),$$

and let two nodes (i, θ_{-i}) and (j, θ'_{-j}) be adjacent if and only if $i \neq j$ and there is a type profile $\hat{\theta}$ satisfying $\hat{\theta}_{-i} = \theta_{-i}$ and $\hat{\theta}_{-j} = \theta'_{-j}$. We refer to G as the *feasibility graph*.

Informally, a node (i, θ_{-i}) is the index for agent i's winning probability when the type profile of the others is θ_{-i}. Two nodes are adjacent if and only if there is a profile $\hat{\theta}$ such that the associated winning probabilities simultaneously appear in the feasibility constraint

$$\sum_{i=1}^{n} \phi_i(\hat{\theta}_{-i}) \leq 1 \tag{1.B.2}$$

of the profile $\hat{\theta}$.

Figure 1.B.1 shows the feasibility graph in an example with two agents; Figure 1.B.2 shows it in an example with three agents.

Given a node $v = (i, \theta_{-i})$ of G, let us write $\phi(v) = \phi_i(\theta_{-i})$. Note that a clique in the feasibility graph is a subset of nodes of V such that the winning probabilities associated with these nodes all appear in the same feasibility constraint (1.B.2). It follows that there is a one-to-one mapping between maximal cliques of G and type profiles. For a DIC mechanism with disposal, the feasibility constraint (1.B.2) may thus be equivalently stated as follows: For all maximal cliques of X of G, we have $\sum_{v \in X} \phi(v) \leq 1$. Thus the set of DIC mechanisms with disposal coincides with the set (1.B.1). One may similarly verify that the set of deterministic DIC mechanisms with disposal coincides with $S(G)$. In view of Theorem 1.24, we deduce:

Lemma 1.25. *All extreme points of the set of DIC mechanisms with disposal are deterministic if and only if G is perfect.*

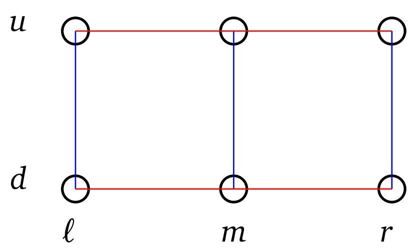

(a) The set of type profiles Θ. Circles represent type profiles.

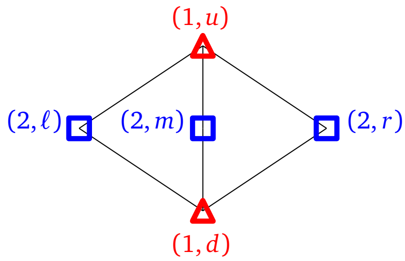

(b) The graph G. Red triangles represent nodes of G that are associated with agent 1. Blue squares represent nodes associated with agent 2.

Figure 1.B.1. The feasibility graph with two agents whose type spaces are $\Theta_1 = \{\ell, m, r\}$ and $\Theta_2 = \{u, d\}$, respectively.

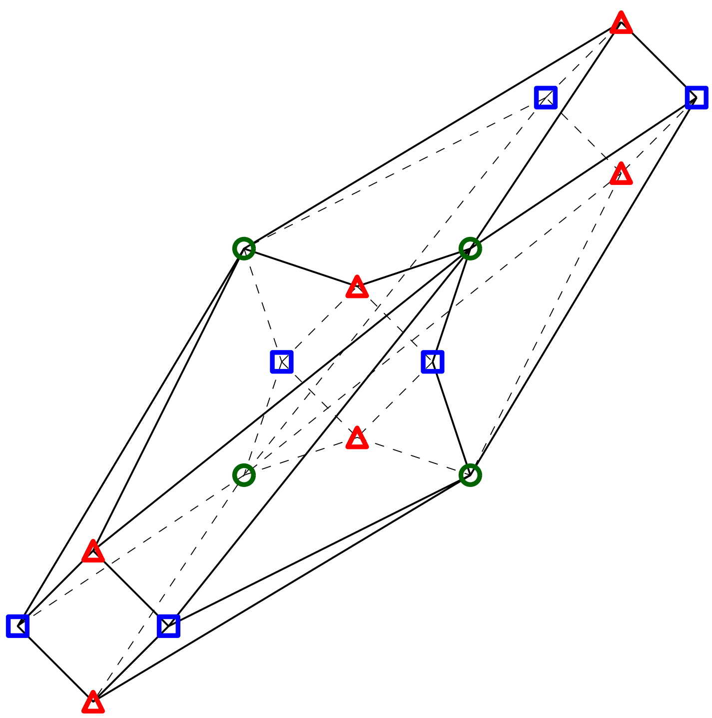

Figure 1.B.2. The feasibility graph G in an example with three agents. Agents 1 and 2 each have two possible types. The nodes of G associated with agents 1 and 2, respectively, are depicted by red triangles and blue squares, respectively. Agent 3 has three possible types; the associated nodes are depicted by green circles. One may view this as the graph G associated with the four-agent environment of Section 1.5.2, except that all nodes of the dummy agent 4 are omitted.

This leads us to the following alternative proof of Theorem 1.22.

Alternative proof of Theorem 1.22. Let $n = 2$. Observe that the node set of G may be partitioned into the sets $\{1\} \times \Theta_2$ and $\{2\} \times \Theta_1$. By definition, two nodes (i, θ_{-i}) and

(j, θ_{-j}) are adjacent only if $i \neq j$. Thus G is bi-partite. Since every bi-partite graph is perfect (Lemma 1.23), the claim follows from Theorem 1.24.

Suppose $|\Theta_i| = 2$ holds for all i. We may relabel the types so that $\Theta_i = \{0, 1\}$ holds for all i. In this case G is the line graph of a bi-partite graph; namely the bi-partite graph with node set $\{0, 1\}^n$ and where two nodes are adjacent if and only if they differ in exactly one entry. The line graph of a bi-partite graph is perfect (Lemma 1.23), and so the claim again follows from Theorem 1.24.

Lastly, suppose $n \geq 3$ and $|\Theta_i| > 2$ for at least one i. We will show that G admits an odd induced cycle of length seven. In view of Lemma 1.23 and Theorem 1.24, this proves that there exists a stochastic extreme point. Let us relabel the agents and types such that the type spaces contain the following subsets of types:

$$\tilde{\Theta}_1 = \{\ell, r\} \quad \text{and} \quad \tilde{\Theta}_2 = \{u, d\} \quad \text{and} \quad \tilde{\Theta}_3 = \{f, c, b\}$$

all hold. Let θ_{-123} be an arbitrary type profile of agents other than 1, 2 and 3 (assuming such agents exist). One may verify that the following is an induced cycle of length seven:

$$\begin{aligned} \big(2, (\ell, c, \theta_{-123})\big) &\longleftrightarrow \big(1, (d, c, \theta_{-123})\big) \\ &\longleftrightarrow \big(3, (r, d, \theta_{-123})\big) \\ &\longleftrightarrow \big(2, (r, b, \theta_{-123})\big) \\ &\longleftrightarrow \big(3, (r, u, \theta_{-123})\big) \\ &\longleftrightarrow \big(1, (u, f, \theta_{-123})\big) \\ &\longleftrightarrow \big(3, (\ell, u, \theta_{-123})\big) \\ \longleftrightarrow \big(2, (\ell, c, \theta_{-123})\big). \end{aligned}$$

□

The proof in the main text for the existence of a stochastic extreme point is slightly more elaborate than the one given above since in the former we explicitly spell out the extreme point. (The proof in the main text uses one of the agents as a dummy, and therefore also works for mechanisms with disposal.) In our view, the advantage of the more elaborate argument is that it facilitates the construction of environments where all deterministic DIC mechanisms fail to be optimal. This lets us give an interpretation as to why it may be optimal to use a lottery. That said, it is clear how the induced cycle defined in the proof of Theorem 1.22 relates to the construction from the main text. The nodes of the cycle correspond to the bold edges of the hyperrectangle in Figure 1.1.

Appendix 1.C Supplementary material: Additional results

This part of the appendix presents results that were previously mentioned in passing.

1.C.1 All extreme points are candidates for optimality

For the following lemma, observe that the set of DIC mechanisms depends only on the number of agents and their type spaces.

Lemma 1.26. *Let $n \in \mathbb{N}$. Let $\Theta_1, \ldots, \Theta_n$ be finite sets, and let $\Theta = \times_{i=1}^n \Theta_i$. If ϕ is an extreme point of the set of DIC mechanisms when there are n agents and the set of type profiles is Θ, then there exists a set Ω of value profiles and a distribution μ over $\Omega \times \Theta$ such that in the environment (n, Ω, Θ, μ) the mechanism ϕ is the unique optimal DIC mechanism.*

Proof of Lemma 1.26. The set of DIC mechanisms is a polytope in Euclidean space (being the set of solutions to a finite system of linear inequalities). Hence all its extreme points are exposed (Aliprantis and Border, 2006, Corollary 7.90). Hence there is a function $p\colon \{1, \ldots, n\} \times \Theta \to \mathbb{R}$ such that for all DIC mechanisms ψ different from ϕ we have $\sum_{i,\theta} p_i(\theta)(\phi_i(\theta) - \psi_i(\theta)) > 0$. By suitably choosing Ω and μ, the function p represents the objective function of our model. For example, one possible choice of Ω and μ is as follows: Let the marginal of μ on Θ be uniform; for all agents i, let Ω_i be the image of p_i; for all θ, conditional on the type profile realizing as θ, let the value of allocating to agent i be $|\Theta| p_i(\theta)$. □

1.C.2 Implementation with deterministic outcome functions

An indirect mechanism specifies a tuple $M = (M_1, \ldots, M_n)$ of finite message sets, and an outcome function $g\colon \times_i M_i \to \Delta\{0, \ldots, n\}$. (Given a finite set X, we denote by ΔX the set of distributions over X.) The outcome function is *deterministic* if for all m the distribution $g(m)$ is degenerate. A *strategy* of agent i in (M, g) is a function $\sigma_i\colon \Theta \to \Delta M_i$; let Σ_i denote the set of strategies of agent i in (M, g). A DIC mechanism ϕ is *implementable (in dominant strategies)* via (M, g) if there is a dominant-strategy equilibrium $(\sigma_1, \ldots, \sigma_n)$ of (M, g) such that all profiles θ satisfy $\phi(\theta) = \sum_m g(m) \prod_i \sigma_i(m_i | \theta_i)$.

Lemma 1.27. *If a stochastic DIC mechanism ϕ is implementable via an indirect mechanism with a deterministic outcome function, then ϕ is not an extreme point of the set of DIC mechanisms.*

Proof of Lemma 1.27. Towards a contradiction, suppose ϕ is an extreme point. As in the proof of Lemma 1.26, we may find $p\colon \{1, \ldots, n\} \times \Theta \to \mathbb{R}$ such that all DIC mechanisms ψ distinct from ϕ satisfy $\sum_{i,\theta} p_i(\theta)(\phi_i(\theta) - \psi_i(\theta)) > 0$. However, since ϕ is implementable via an indirect mechanism with a deterministic outcome function, Proposition 1 of Jarman and Meisner (2017) implies that there is a deterministic DIC mechanism ψ such that

$$\forall_{\theta \in \Theta}, \quad \sum_i p_i(\theta)(\phi_i(\theta) - \psi_i(\theta)) \leq 0.$$

Hence $\sum_{i,\theta} p_i(\theta)(\phi_i(\theta) - \psi_i(\theta)) \leq 0$. Since ϕ is stochastic, we have $\psi \neq \phi$; contradiction. $\square$

1.C.3 Total unimodularity

This section of the appendix discusses another potential approach for showing that all extreme points are deterministic. Our aim is to explain why this approach does not help us for the proof of Theorem 1.5 in the difficult case with three agents.

For a function $\phi : \Theta \to [0,1]^n$ to be a DIC mechanism, the function should satisfy the following:

$$\begin{aligned} \forall_{i,\theta}, \quad 1 &\geq \phi_i(\theta) \\ \forall_{i,\theta_i,\theta_i',\theta_{-i}}, \quad 0 &\geq \phi_i(\theta_i, \theta_{-i}) - \phi_i(\theta_i', \theta_{-i}) \geq 0 \\ \forall_\theta, \quad 1 &\geq \sum_i \phi_i(\theta) \geq 1 \end{aligned} \tag{1.C.1}$$

For a suitable matrix A and vector b, the set of DIC mechanisms is then the polytope $\{\phi : A\phi \geq b, \phi \geq 0\}$. Here, the matrix A has one row for every constraint in (1.C.1) (after splitting the constraints into one-sided inequalities). Each column of A identifies a pair of the form (i, θ).

A matrix or a vector is *integral* if its entries are all in $\mathbb{Z}$. A polytope is *integral* if all its extreme points are integral. In this language, all extreme points of the set of DIC mechanisms are deterministic if and only if the polytope $\{\phi : A\phi \geq b, \phi \geq 0\}$ is integral.

Recall that a matrix is *totally unimodular* if all its square submatrices have a determinant equal to -1, 0, or 1. A submatrix of a totally unimodular matrix is itself totally unimodular.

Our interest in total unimodularity is due the Hoffman-Kruskal theorem (Korte and Vygen, 2018, Theorem 5.21).

Theorem 1.28. *An integral matrix A is totally unimodular if and only if for all integral vectors b all extreme points of the set $\{\phi : A\phi \geq b, \phi \geq 0\}$ are integral.*

Thus a sufficient condition for all extreme points of the set of DIC mechanisms to be deterministic is that the constraint matrix A be totally unimodular. Unfortunately:

Lemma 1.29. *For all agents i, let $|\Theta_i| \geq 2$. Let $n = 3$. If there exists i such that $|\Theta_i| \geq 3$, then A is not totally unimodular.*

Proof of Lemma 1.29. Towards a contradiction, suppose A is totally unimodular. Consider the constraint matrix $\tilde{A}$ and vector $\tilde{b}$ that define the set of DIC mechanisms with disposal (where such mechanisms are defined in Appendix 1.B). That is, ϕ is a DIC mechanism with disposal if and only if $\tilde{A}\phi \geq \tilde{b}$ and $\phi \geq 0$. Notice that $\tilde{A}$ is obtained from A by dropping all rows corresponding to constraints of the form $\sum_i \phi_i(\theta) \geq 1$; the vector $\tilde{b}$ is obtained from b by dropping the corresponding entries. In particular, the matrix $\tilde{A}$ is a submatrix of A. Since A is totally unimodular, we conclude that $\tilde{A}$ is totally unimodular. We infer from Theorem 1.28 that all extreme points of the set of DIC mechanism with disposal are deterministic. Since $n = 3$, since all agents have at least binary types, and since at least one agent has non-binary types, we have a contradiction to Theorem 1.22. $\square$

We can give an alternative proof of Lemma 1.29 that does not require Theorem 1.22. Consider the following characterization of total unimodularity due to Ghouila-Houri (1962) (Korte and Vygen, 2018, Theorem 5.25).

Theorem 1.30. *A matrix A with entries in $\{-1, 0, 1\}$ is totally unimodular if and only if all subsets C of columns of A satisfy the following: There exists a partition of C into subsets C^+ and C^- such that for all rows r of A we have*

$$\left(\sum_{c \in C^+} A(r,c) - \sum_{c \in C^-} A(r,c) \right) \in \{-1, 0, 1\}. \tag{1.C.2}$$

Alternative proof of Lemma 1.29. Let us relabel the agents and types such that the type spaces contain the following subsets:

$$\tilde{\Theta}_1 = \{\ell, r\} \quad \text{and} \quad \tilde{\Theta}_2 = \{u, d\} \quad \text{and} \quad \tilde{\Theta}_3 = \{f, c, b\}$$

Fixing an arbitrary type profile θ_{-123} of agents other than 1, 2, and 3, let us define the type profiles $\{\theta^a, \theta^b, \theta^c, \theta^e, \theta^f, \theta^g\}$ as in (1.2) in Section 1.5.2. That is, let

$$\begin{aligned} \theta^a = (\ell, d, c, \theta_{-123}), \quad \theta^b &= (r, d, c, \theta_{-123}), \quad \theta^c = (r, d, b, \theta_{-123}), \\ \theta^d &= (r, u, b, \theta_{-123}), \quad \theta^e = (r, u, f, \theta_{-123}), \\ \theta^f &= (\ell, u, f, \theta_{-123}), \quad \theta^g = (\ell, u, c, \theta_{-123}). \end{aligned}$$

Recall that each column of A corresponds to an entry of the form (i, θ). We will argue that the following set C of columns does not admit a partition in the sense of Theorem 1.30.

$$\begin{aligned} C = \{&(1, \theta^a), (1, \theta^b), (3, \theta^b), (3, \theta^c), \\ &(2, \theta^c), (2, \theta^d), (3, \theta^d), (3, \theta^e), \\ &(1, \theta^e), (1, \theta^f), (3, \theta^f), (3, \theta^g), \\ &(2, \theta^g), (2, \theta^a)\} \end{aligned}$$

Towards a contradiction, suppose C does admit a partition into sets C^+ and C^- in the sense of Theorem 1.30. Let us assume $(1, \theta^a) \in C^+$, the other case being similar. Note that θ^a and θ^b differ only in the type of agent 1. Consider the row of A corresponding to the DIC constraint for agent 1 at these type profiles. By referring to (1.C.2) for this row, we deduce $(1, \theta^b) \in C^+$. Next, via a similar argument, the constraint that the object is allocated at θ^b requires $(3, \theta^b) \in C^-$. Continuing in this manner, it is easy to see that $(1, \theta^a)$ must be in C^-. Since $(1, \theta^a)$ is assumed to be in C^+, we have a contradiction to the assumption that C^+ and C^- are a partition of C. $\square$

1.C.4 Maximum weight perfect hypergraph matching

In this section, we explain that the problem of finding an optimal deterministic DIC mechanism corresponds to finding a maximum weight perfect matching on a certain hypergraph.

The hypergraph has as vertices the set of type profiles. Its hyperedges are those type profiles along which the type of exactly one agent i varies across Θ_i. That is, a set of type profiles e is a hyperedge if and only if there exist $i \in \{1, \ldots, n\}$ and $\theta_{-i} \in \Theta_{-i}$ such that $e = \{(\theta_i, \theta_{-i}) : \theta_i \in \Theta_i\}$. We index this hyperedge by (i, θ_{-i}). The weight attached to hyperedge (i, θ_{-i}) is $\mathbb{E}_{\omega_i}[\omega_i | \theta_{-i}]$.

In a matching of this hypergraph, including edge (i, θ_{-i}) in the matching corresponds to allocating to agent i at all type profiles incident to (i, θ_{-i}); this respects DIC for agent i. In a perfect matching, each type profile is covered by some edge; this respects the requirement that the object is always allocated.

If we relax the requirement that the object is always allocated (Appendix 1.B), we instead consider the larger set of all matchings on the hypergraph. Such a matching can also be interpreted as a stable set of the feasibility graph introduced in Appendix 1.B.2.2.

References

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias. 2012. "Targeting the poor: evidence from a field experiment in Indonesia." *American Economic Review* 102 (4): 1206–40. [22]

Aliprantis, Charalambos D., and Kim C. Border. 2006. *Infinite Dimensional Analysis : A Hitchhikers Guide.* Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-29587-9. [10, 46]

Alon, Noga, Felix Fischer, Ariel Procaccia, and Moshe Tennenholtz. 2011. "Sum of us: Strategyproof selection from the selectors." In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge,* 101–10. [4, 6, 7, 21]

Aziz, Haris, Omer Lev, Nicholas Mattei, Jeffrey Rosenschein, and Toby Walsh. 2016. "Strategyproof peer selection: Mechanisms, analyses, and experiments." In *Proceedings of the AAAI Conference on Artificial Intelligence,* vol. 30. 1. [6]

Aziz, Haris, Omer Lev, Nicholas Mattei, Jeffrey S Rosenschein, and Toby Walsh. 2019. "Strategyproof peer selection using randomization, partitioning, and apportionment." *Artificial Intelligence* 275: 295–309. [6]

Baumann, Leonie. 2018. "Self-ratings and peer review." https://sites.google.com/site/leoniebaumann/. [7]

Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman. 2014. "Optimal allocation with costly verification." *American Economic Review* 104 (12): 3779–813. [4, 6]

Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman. 2019. "Mechanisms with evidence: Commitment and robustness." *Econometrica* 87 (2): 529–66. [4]

Bhaskar, Dhruva, and Evan Sadler. 2019. "Resource Allocation with Positive Externalities." [6]

Bjelde, Antje, Felix Fischer, and Max Klimm. 2017. "Impartial selection and the power of up to two choices." *ACM Transactions on Economics and Computation (TEAC)* 5 (4): 1–20. [6]

Bloch, Francis, Bhaskar Dutta, and Marcin Dziubiński. 2022. "Selecting a Winner with Impartial Referees." https://www.sites.google.com/site/francisbloch1/papers. [7]

Bloch, Francis, and Matthew Olckers. 2021. "Friend-based ranking in practice." *AEA Papers and Proceedings* 111: 567–71. [7]

Bloch, Francis, and Matthew Olckers. 2022. "Friend-based ranking." *American Economic Journal: Microeconomics* 14 (2): 176–214. [7]

Bousquet, Nicolas, Sergey Norin, and Adrian Vetta. 2014. "A near-optimal mechanism for impartial selection." In *International Conference on Web and Internet Economics,* 133–46. Springer. [6]

Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. "Designing random allocation mechanisms: Theory and applications." *American Economic Review* 103 (2): 585–623. [7]

Caragiannis, Ioannis, George Christodoulou, and Nicos Protopapas. 2019. "Impartial selection with additive approximation guarantees." In *International Symposium on Algorithmic Game Theory,* 269–83. Springer. [6]

Caragiannis, Ioannis, George Christodoulou, and Nicos Protopapas. 2021. "Impartial selection with prior information." https://doi.org/https://doi.org/10.48550/arXiv.2102.09002. [6]

Chakravarty, Surajeet, and Todd R. Kaplan. 2013. "Optimal allocation without transfer payments." *Games and Economic Behavior* 77 (1): 1–20. [6]

Chen, Yi-Chun, Wei He, Jiangtao Li, and Yeneng Sun. 2019. "Equivalence of Stochastic and Deterministic Mechanisms." *Econometrica* 87 (4): 1367–90. [4, 7, 13, 16]

Chvátal, Vašek. 1975. "On certain polytopes associated with graphs." *Journal of Combinatorial Theory, Series B* 18 (2): 138–54. [7, 17, 42]

Condorelli, Daniele. 2012. “What money can’t buy: Efficient mechanism design with costly types.” *Games and Economic Behavior* 75 (2): 613–24. [6]

Crémer, Jacques, and Richard P. McLean. 1985. “Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent.” *Econometrica,* 345–61. [7]

Crémer, Jacques, and Richard P. McLean. 1988. “Full extraction of the surplus in Bayesian and dominant strategy auctions.” *Econometrica,* 1247–57. [7]

de Clippel, Geoffroy, Kfir Eliaz, Daniel Fershtman, and Kareen Rozen. 2021. “On selecting the right agent.” *Theoretical Economics* 16 (2): 381–402. [21]

de Clippel, Geoffroy, Hervé Moulin, and Nicolaus Tideman. 2008. “Impartial division of a dollar.” *Journal of Economic Theory* 139 (1): 176–91. [6]

Edelman, Paul H, and Attila Por. 2021. “A new axiomatic approach to the impartial nomination problem.” *Games and Economic Behavior* 130: 443–51. [6]

Epitropou, Markos, and Rakesh Vohra. 2019. “Optimal On-Line Allocation Rules with Verification.” In *International Symposium on Algorithmic Game Theory,* 3–17. Springer. [6]

Erlanson, Albin, and Andreas Kleiner. 2019. “A note on optimal allocation with costly verification.” *Journal of Mathematical Economics* 84: 56–62. [6]

Fischer, Felix, and Max Klimm. 2015. “Optimal impartial selection.” *SIAM Journal on Computing* 44 (5): 1263–85. [6]

Ghouila-Houri, Alain. 1962. “Caractérisation des matrices totalement unimodulaires.” *Comptes Redus Hebdomadaires des Séances de l’Académie des Sciences (Paris)* 254: 1192–94. [48]

Goldlücke, Susanne, and Thomas Tröger. 2020. “The multiple-volunteers principle.” *Available at SSRN 3753985.* [6]

Guo, Yingni, and Johannes Hörner. 2021. “Dynamic allocation without money.” https://hal.archives-ouvertes.fr/hal-03187506. [6, 21]

Holzman, Ron, and Hervé Moulin. 2013. “Impartial nominations for a prize.” *Econometrica* 81 (1): 173–96. Publisher: Wiley Online Library. [4–7, 9, 20–22]

Jarman, Felix, and Vincent Meisner. 2017. “Deterministic mechanisms, the revelation principle, and ex-post constraints.” *Economics Letters* 161: 96–98. [7, 46]

Kato, Miki, and Shinji Ohseto. 2002. “Toward general impossibility theorems in pure exchange economies.” *Social Choice and Welfare* 19 (3): 659–64. [9]

Kattwinkel, Deniz. 2019. “Allocation with Correlated Information: Too good to be true.” https://sites.google.com/view/kattwinkel. [7]

Kattwinkel, Deniz, and Jan Knoepfle. 2021. “Costless Information and Costly Verification: A Case for Transparency.” https://dx.doi.org/10.2139/ssrn.3426817. [7]

Kattwinkel, Deniz, Axel Niemeyer, Justus Preusser, and Alexander Winter. 2022. “Mechanisms without transfers for fully biased agents.” https://arxiv.org/abs/2205.10910. [7]

Korte, Bernhard, and Jens Vygen. 2018. *Combinatorial Optimization: Theory and Algorithms.* 6th ed. 2018. Algorithms and Combinatorics; 21. Berlin, Heidelberg: Springer Berlin Heidelberg. [12, 31, 42, 47, 48]

Lev, Omer, Nicholas Mattei, Paolo Turrini, and Stanislav Zhydkov. 2021. “Peer Selection with Noisy Assessments.” https://doi.org/https://doi.org/10.48550/arXiv.2107.10121. [6]

Li, Yunan. 2020. “Mechanism design with costly verification and limited punishments.” *Journal of Economic Theory* 186: 105000. [6]

Lipnowski, Elliot, and Joao Ramos. 2020. “Repeated delegation.” *Journal of Economic Theory* 188: 105040. [21]

Mackenzie, Andrew. 2015. “Symmetry and impartial lotteries.” *Games and Economic Behavior* 94: 15–28. https://doi.org/10.1016/j.geb.2015.08.007. [5, 6, 20, 21]

Mackenzie, Andrew. 2020. “An axiomatic analysis of the papal conclave.” *Economic Theory* 69 (3): 713–43. [6, 20, 21]

Mattei, Nicholas, Paolo Turrini, and Stanislav Zhydkov. 2020. “PeerNomination: Relaxing exactness for increased accuracy in peer selection.” https://doi.org/https://doi.org/10.48550/arXiv.2004.14939. [6]

McAfee, R. Preston, and Philip J. Reny. 1992. “Correlated information and mechanism design.” *Econometrica,* 395–421. [7]

Mylovanov, Tymofiy, and Andriy Zapechelnyuk. 2017. “Optimal allocation with ex post verification and limited penalties.” *American Economic Review* 107 (9): 2666–94. [6]

Olckers, Matthew, and Toby Walsh. 2022. “Manipulation and Peer Mechanisms: A Survey.” *arXiv preprint arXiv:2210.01984.* [6]

Pycia, Marek, and M. Utku Ünver. 2015. “Decomposing random mechanisms.” *Journal of Mathematical Economics* 61: 21–33. [4, 7, 12, 13]

Rivera Mora, Ernesto. 2022. “Deterministic Mechanism Design.” https://www.ernestoriveramora.com/research. [7, 13]

Tamura, Shohei. 2016. “Characterizing minimal impartial rules for awarding prizes.” *Games and Economic Behavior* 95: 41–46. Publisher: Elsevier. [6]

Tamura, Shohei, and Shinji Ohseto. 2014. “Impartial nomination correspondences.” *Social Choice and Welfare* 43 (1): 47–54. [6]

Chapter 2

Mechanisms Without Transfers for Fully Biased Agents

2.1 Introduction

A principal has to decide between two options. Which one she prefers depends on the private information of two agents. One agent always prefers the first option; the other always prefers the second. Transfers are infeasible. The principal designs and commits to a mechanism: a mapping from reported information profiles to a – potentially randomized – decision. One prominent example of such a setting is the allocation of a fixed amount of money:

Example 2.1 (Budget allocation)**.** Upper management has endowed a division manager with a fixed budget. The manager can divide these funds between her two departments L, R. Her objective is to maximize expected returns. Department heads $i = \ell, r$ hold private information θ_i about the future marginal returns y_L, y_R and want to maximize their department's budget. Formally, $(\theta_\ell, \theta_r, y_L, y_R)$ follows some joint distribution and conditional on the private information the manager's marginal return from allocating \$1 to L instead of R is $v(\theta_\ell, \theta_r) = E[y_L - y_R | \theta_\ell, \theta_r]$.

In this setting, we characterize all implementable mechanisms without transfers under arbitrary correlation. We find a connection between our mechanism design setting and a zero-sum game. Incentive compatibility of a mechanism given a type distribution corresponds to this distribution being a correlated equilibrium in the game induced by the mechanism.

Crémer and McLean's (1985; 1988) results for the corresponding setting with transfers suggest that the principal should be able to exploit correlation to induce truthful reporting. We define a preorder on type distributions and find that correlation has the opposite effect in our setting: it restricts the set of implementable mechanisms. Under Crémer and McLean's full-rank condition on the the joint type distribution, the set of implementable mechanisms collapses to the mechanisms that ignore the agents' reports entirely. In particular, under the full-rank condition, the principal cannot do better than choosing her ex-ante preferred option.

We then give necessary and sufficient conditions (on the joint type distribution) for the existence of a "profitable" mechanism that allows the principal to do better than her ex-ante preferred option. When she is ex-ante indifferent between the two options, the existence of a profitable mechanism is equivalent to a non-additive payoff structure. Informally speaking, in the money-allocation example, non-additivity means that one department has valuable information about the expected marginal return of the *other* department.

When the principal is not ex-ante indifferent, a key insight is that choosing a mechanism corresponds to introducing endogenous correlation. The existence of a profitable mechanism depends on the value of a related optimal transport problem in which the principal chooses this endogenous correlation structure. Incentive constraints translate into an equal marginals condition and an orthogonality constraint between the exogenously given type distribution and the endogenously chosen one.

One application of our results is the problem of allocating a single nondisposable good between two agents. In section 6, we extend our setting and study the problem of allocating a (potentially disposable) good among multiple agents under independence. When the good has to be allocated, we find that a profitable mechanism exists if and only if a generalized version of the additivity condition is violated. Under free disposal, a profitable mechanism exists if and only if there is an agent such that the principal's value from allocating to that agent depends on the types of other agents.

More broadly, our results convey that there is large class of settings without transfers where the principal can profit from designing a mechanism that elicits the agent's information despite their opposed interests. This scope for communication does not rely on any correlation of the agents' information but instead on types interdependence in the principal's preferences.

2.2 Model

There is a principal (she), two agents $i = \ell, r$ and a decision: L or R. Agent ℓ (he) always prefers L; agent r (he) always prefers R. Agents enjoy utility 1 if their favored decision is taken and 0 otherwise.[1] Each agent has a private *type* $\theta_i \in \Theta_i$ ($|\Theta_i| < \infty$) and the type profile $\theta = (\theta_\ell, \theta_r)$ is drawn from a commonly known distribution $\pi(\theta_\ell, \theta_r)$ with positive[2] marginals π_ℓ, π_r. Let Π be the set of joint type distributions with positive marginals and let $\Pi(\pi_\ell, \pi_r)$ be the set of joint type dis-

1. All of our results would apply unchanged if agents receive utility $\bar{u}_i(\theta_i)$ from their preferred decision and utility $\underline{u}_i(\theta_i)$ from their less preferred decision, where $\bar{u}_i > \underline{u}_i$.

2. This is without loss. Note that we do not assume full support.

tributions with marginals π_i. The independent type distribution with marginals π_ℓ and π_r is denoted by $\pi_\ell \pi_r$.

The principal designs and commits to a mechanism. By the revelation principle she can restrict attention to direct, (Bayesian) incentive-compatible mechanisms $x\colon \Theta = \Theta_\ell \times \Theta_r \to [0,1]$, where $x(\theta_\ell, \theta_r)$ denotes the probability that L is chosen if agent ℓ reports θ_ℓ and agent r reports θ_r. From now on we refer to direct mechanisms simply as mechanisms.

If the realized type profile is θ then the principal receives a payoff of $v_L(\theta)$ from L and of $v_R(\theta)$ from R. Without loss of generality, her ex-ante preferred option is R, and v_R is normalized to 0; that is, $E_\pi[v_L(\boldsymbol{\theta})] \le 0 = v_R$.[3] From now on we denote $v_L = v$.

The principal's problem then reads:

$$\begin{aligned}
&\max_{0 \le x \le 1} && E_\pi[v(\boldsymbol{\theta})x(\boldsymbol{\theta})] && \\
&\text{s.t.} && E_\pi[x(\theta_\ell, \boldsymbol{\theta}_r)|\theta_\ell] \ge E_\pi[x(\theta'_\ell, \boldsymbol{\theta}_r)|\theta_\ell] \quad \forall \theta_\ell, \theta'_\ell && (IC_\ell)\\
& && E_\pi[x(\boldsymbol{\theta}_\ell, \theta_r)|\theta_r] \le E_\pi[x(\boldsymbol{\theta}_\ell, \theta'_r)|\theta_\ell] \quad \forall \theta_r, \theta'_r && (IC_r)
\end{aligned}$$

Given $\pi \in \Pi$, let the set of IC mechanisms be $\mathscr{X}(\pi)$. A mechanism is said to be *profitable* if it is IC and yields the principal a strictly greater payoff than choosing her ex-ante preferred option R without consulting the agents. Given our normalization of the principal's payoff, an IC mechanism x is profitable if and only if $E_\pi[v(\boldsymbol{\theta})x(\boldsymbol{\theta})] > 0$.

2.3 Implementation

In this section we characterize the set of IC mechanisms given a type-distribution. The proof is based on the observation that incentive-compatibility can be phrased in terms of the correlated equilibria of an auxiliary two-player zero-sum game, as we explain next.

Let $\pi \in \Pi$ and let $x \in \mathscr{X}(\pi)$ be an IC mechanism. The IC conditions read

$$\sum_{\theta_r} \pi(\theta_r|\theta_\ell)x(\theta_\ell, \theta_r) \ge \sum_{\theta_r} \pi(\theta_r|\theta_\ell)x(\theta'_\ell, \theta_r) \quad \forall\, \theta_\ell, \theta'_\ell \qquad (IC'_\ell)$$

$$\sum_{\theta_\ell} \pi(\theta_\ell|\theta_r)x(\theta_\ell, \theta_r) \le \sum_{\theta_\ell} \pi(\theta_\ell|\theta_r)x(\theta_\ell, \theta'_r) \quad \forall\, \theta_r, \theta'_r. \qquad (IC'_r)$$

Consider now the auxiliary two-player zero-sum game G in which the Maximizer chooses θ_ℓ, the Minimizer chooses θ_r and the objective (i.e. the Maximizer's payoff if θ_ℓ and θ_r is chosen) is $x(\theta_\ell, \theta_r)$. In this game we can interpret π as a *correlated strategy* under which the Maximizer's payoff is $\sum_{\theta_\ell} \sum_{\theta_r} \pi(\theta_\ell, \theta_r)x(\theta_\ell, \theta_r)$. With this

3. Bold symbols denote random variables.

interpretation the IC conditions become *obedience conditions* and π becomes a *correlated equilibrium* of G:

Lemma 2.1. *A mechanism x is IC under some type distribution $\pi \in \Pi$ if and only if π is a correlated equilibrium of the auxiliary two-player zero-sum game in which the Maximizer chooses $\theta_\ell \in \Theta_\ell$, the Minimizer chooses $\theta_r \in \Theta_r$ and the Maximizer's payoff is $x(\theta_\ell, \theta_r)$.*

Note that under the mechanism design interpretation π is an exogenous part of the environment while x is endogenous. In the auxiliary game the roles are exactly flipped: x is an exogenous while π is endogenous.

Proposition 2.2. *Let $\pi \in \Pi$ and let x be some mechanism. The following are equivalent.*

(i) x is IC under π.

(ii) For each agent i, each type θ_i, and each report θ_i', we have $E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})|\theta_i] = E_\pi[x(\boldsymbol{\theta})]$.

Moreover, if x is IC then $E_\pi[x(\boldsymbol{\theta})] = \bar{x}$, where

$$\bar{x} = \max_{\sigma_\ell \in \Delta\Theta_\ell} \min_{\sigma_r \in \Delta\Theta_r} \sum_{\theta_\ell} \sum_{\theta_r} \sigma_\ell(\theta_\ell)\sigma_r(\theta_r)x(\theta_\ell, \theta_r)$$

is the maximin value of the auxiliary game.

Proposition 2.2 says that a mechanism is IC if and only if each type of each agent is indifferent between every possible report and each type's expectations of x are given by the distribution-independent constant $\bar{x}$.

Proof of Proposition 2.2. Any mechanism that satisfies (ii) is clearly IC. To show the converse let $\pi \in \Pi$ and let x be an IC mechanism. By Lemma 2.1, π is a correlated equilibrium of the auxiliary game G in which the Maximizer chooses θ_ℓ, the Minimizer chooses θ_r and the Maximizer's payoff from such an action profile is $x(\theta_\ell, \theta_r)$.

Suppose now that the Minimizer obeys his recommendation under the correlated equilibrium π. Suppose further that the Maximizer disobeys his recommendation under π and instead plays the mixed strategy π_ℓ. As disobediance is not profitable, we get

$$\sum_{\theta_\ell} \sum_{\theta_r} \pi(\theta_\ell, \theta_r)x(\theta_\ell, \theta_r) \geq \sum_{\theta_\ell} \sum_{\theta_r} \pi(\theta_\ell, \theta_r) \sum_{\theta_\ell'} \pi_\ell(\theta_\ell')x(\theta_\ell', \theta_r).$$

But the last term is simply $\sum_{\theta_\ell} \sum_{\theta_r} \pi_\ell(\theta_\ell)\pi_r(\theta_r)x(\theta_\ell, \theta_r)$ and so we obtain

$$\sum_{\theta_\ell} \sum_{\theta_r} \pi(\theta_\ell, \theta_r)x(\theta_\ell, \theta_r) \geq \sum_{\theta_\ell} \sum_{\theta_r} \pi_\ell(\theta_\ell)\pi_r(\theta_r)x(\theta_\ell, \theta_r). \tag{2.1}$$

The symmetric argument for the Minimizer implies that the opposite inequality to (2.1) must also hold. We conclude that (2.1) must hold with equality. Finally, since π_ℓ has full support if there were some pair $\theta_\ell, \theta'_\ell$ for which the inequality in the obedience constraint (IC'_ℓ) were strict then (2.1) could not hold with equality. Hence (IC'_ℓ) must always bind. A similar argument shows that (IC'_r) must always bind. Thus, for all i, if player i is recommended some action θ_i then he is indifferent between all actions and his interim expectation of x is $\bar{x}_i(\theta_i) = E_\pi[x(\boldsymbol{\theta})|\theta_i]$. We will now show that the interim expectations $\bar{x}_i(\theta_i)$ are actually all the same.

Let $\sigma = (\sigma_\ell, \sigma_r)$ be a Nash equilibrium of G. (Existence follows from finiteness of G.) Let $\bar{x} = \sum_{\theta_\ell} \sigma_\ell(\theta_\ell) \sum_{\theta_r} \sigma_r(\theta_r) x(\theta_\ell, \theta_r)$ be the Maximizer's expected payoff under σ. Then for any θ_r it holds that

$$\begin{aligned} \bar{x} &\geq \sum_{\theta_\ell} \pi(\theta_\ell|\theta_r) \sum_{\tilde{\theta}_r} \sigma_r(\tilde{\theta}_r) x(\theta_\ell, \tilde{\theta}_r) \\ &= \sum_{\tilde{\theta}_r} \sigma_r(\tilde{\theta}_r) \sum_{\theta_\ell} \pi(\theta_\ell|\theta_r) x(\theta_\ell, \tilde{\theta}_r) \\ &\geq \sum_{\theta_\ell} \pi(\theta_\ell|\theta_r) x(\theta_\ell, \theta_r) \; = \; \bar{x}_r(\theta_r), \end{aligned}$$

where the first inequality holds since the mixed strategy $\pi(\cdot|\theta_r)$ is not a profitable deviation from σ_ℓ; the second inequality follows from type θ_r's IC constraint for $\tilde{\theta}_r$ and the last equality is by definition. Combining this inequality with the corresponding inequality for the other player we thus have that

$$\bar{x}_r(\theta_r) \; \leq \; \bar{x} \; \leq \; \bar{x}_\ell(\theta_\ell) \quad \forall \theta_\ell, \theta_r.$$

Since the terms on the left and the right hand side of the above inequalities are equal in expectation and all θ_ℓ and θ_r occur with positive probability both inequalities above must always bind. That is to say, for each agent i, $E_\pi[x(\theta'_i, \boldsymbol{\theta}_{-i})|\theta_i] = \bar{x}$ for any θ_i and θ'_i. Finally, note that $\bar{x} = \max_{\sigma_\ell \in \Delta\Theta_\ell} \min_{\sigma_r \in \Delta\Theta_r} \sum_{\theta_\ell} \sum_{\theta_r} \sigma_\ell(\theta_\ell) \sigma_r(\theta_r) x(\theta_\ell, \theta_r)$ holds since (σ_ℓ, σ_r) is a Nash equilibrium of the zero sum game G. □

2.4 Comparative statics for implementation

In this section, we study how the set of implementable mechanisms depends on the type distribution. We define a preorder on distributions and derive a monotone comparative statics result for the correspondence $\pi \mapsto \mathscr{X}(\pi)$. We conclude that correlation has a restrictive effect.

Definition 2.1. Let $\tau^0, \tau^1, \ldots, \tau^k \in \Delta\Theta_{-i}$ be a collection of beliefs over types of agent $-i$. Then $\{\tau^1, \ldots, \tau^k\}$ is said to *span* τ^0 if there exist coefficients $\alpha_1, \ldots, \alpha_k \in \mathbb{R}$ such that

$$\tau^0(\theta_{-i}) \; = \; \sum_{j=1}^{k} \tau^j(\theta_{-i}) \alpha_j \quad \forall \theta_{-i}.$$

Given joint type distributions $\pi, \tilde{\pi} \in \Pi$, we say π *spans* $\tilde{\pi}$ if for all $i = \ell, r$ and θ_i the collection $\{\pi(\cdot|\tilde{\theta}_i) : \tilde{\theta}_i \in \Theta_i\}$ spans $\tilde{\pi}(\cdot|\theta_i)$.

Hence π spans $\tilde{\pi}$ if each interim belief an agent can hold under $\tilde{\pi}$ is a linear combination of some interim beliefs that he can hold under π. Spanning is reflexive and transitive but not anti-symmetric and therefore defines a preorder.

Example 2.2. Let $\pi \in \Pi$ with marginals π_i. Then π spans the independent type distribution $\tilde{\pi} = \pi_\ell \pi_r$ since

$$\pi_i(\theta_i) = \sum_{\theta_{-i}} \pi(\theta_i|\theta_{-i})\pi_{-i}(\theta_{-i}) \quad \forall \theta_i \forall i.$$

Example 2.3. A joint distribution $\pi \in \Pi$ spans every other joint distribution $\tilde{\pi} \in \Pi$ if and only if the matrix $(\pi(\theta_\ell, \theta_r)) \in \mathbb{R}^{\Theta_\ell \times \Theta_r}$ has full column-rank and full row-rank. This is exactly the condition introduced by Crémer and McLean (see Assumption 4 in their 1985 paper and Theorem 1 in their 1988 paper).

Our first application of the spanning relation shows that the set of IC mechanisms cannot shrink when passing from π to some other type distribution $\tilde{\pi}$ that is spanned by π.

Proposition 2.3. *Let $\pi, \tilde{\pi} \in \Pi$ be type distributions. If π spans $\tilde{\pi}$ then*

$$\mathscr{X}(\pi) \subset \mathscr{X}(\tilde{\pi}).$$

Proof. By Proposition 2.2 a mechanism x is IC under π if and only if

$$\sum_{\theta_{-i}} \pi(\theta_{-i}|\theta_i)(x(\theta_i', \theta_{-i}) - \bar{x}) = 0 \quad \forall \theta_i, \theta_i', i = \ell, r.$$

Now let x be IC under π and consider some $\tilde{\pi} \in \Pi$ spanned by π. By definition, there exist coefficients $\alpha_i(\theta_i, \tilde{\theta}_i)$ such that

$$\tilde{\pi}(\theta_{-i}|\theta_i) = \sum_{\tilde{\theta}_i} \pi(\theta_{-i}|\tilde{\theta}_i)\alpha_i(\theta_i, \tilde{\theta}_i) \quad \forall \theta_i, \theta_{-i}, i = \ell, r.$$

But then x must also be IC under $\tilde{\pi}$ because for all θ_i, θ_i':

$$\sum_{\theta_{-i}} \tilde{\pi}(\theta_{-i}|\theta_i)(x(\theta_i', \theta_{-i}) - \bar{x}) = \sum_{\tilde{\theta}_i} \alpha_i(\theta_i, \tilde{\theta}_i) \sum_{\theta_{-i}} \pi(\theta_{-i}|\tilde{\theta}_i)(x(\theta_i', \theta_{-i}) - \bar{x}) = 0.$$

□

The proof (appendix) of the next result is another application of the spanning relation.

Proposition 2.4. *Let $\pi \in \Pi$ with marginals π_i and let x be some mechanism. Then:*

(1) If x is IC under π then x is also IC under the independent type distribution $\tilde{\pi} = \pi_\ell \pi_r$.

(2) If the matrix $(\pi(\theta_\ell, \theta_r)) \in \mathbb{R}^{\Theta_\ell \times \Theta_r}$ has full rank then only constant mechanisms are IC.

The maximal elements of the spanning preorder are exactly the full-rank distributions and its minimal elements are exactly the independent distributions.

Crémer and McLean (1985) show in a setting with transfers that full rank correlation makes it possible to implement any allocation rule while paying zero information rents.[4] We show that under the same full-rank condition only mechanisms that ignore the agents' reports are IC. Absent full rank correlation, any mechanism that is IC under correlation must also be IC when types are independent. This shows that the spirit of Crémer and McLean's results is inverted in our setting. The next example illustrates this difference.

Example 2.4. Assume $\Theta_\ell = \Theta_r = \{-1, 1\}$, $\pi_\ell = \pi_r = \frac{1}{2}$ and $v(\theta) = \theta_\ell \theta_r$. Both options yield the principal an ex-ante expected payoff of 0 while the first best mechanism x^* would choose L iff $\theta_\ell = \theta_r$ and yield $E[v(\boldsymbol{\theta})x^*(\boldsymbol{\theta})] = \frac{1}{2} > 0$. If types are independent then x^* is actually IC because from each agent's perspective, any report will lead to the same probability of L. Now assume instead that types are correlated and that π is given by

$$\begin{array}{c|cc|} & -1 & 1 \\ \cline{2-3} -1 & 0.25-\varepsilon & 0.25+\varepsilon \\ 1 & 0.25+\varepsilon & 0.25-\varepsilon \\ \cline{2-3} \end{array}$$

where $0 < \varepsilon \leq \frac{1}{4}$ is arbitrary. Then x^* is not IC anymore: For example, type 1 of agent ℓ would infer from his type that the other agent's type is probably -1 and would therefore claim to be type -1 instead of being truthful. In fact, since the distribution matrix has full rank, Proposition 2.4 implies that the only remaining IC mechanisms are constant.

2.5 Profitable mechanisms

In this section we investigate when the principal can design a profitable mechanism. We attack this question from two different angles. Our first characterization is in

4. The full-rank condition is often seen as generic. In many applications, though, it is not satisfied even when types are correlated. For example, assume that there exists a finite underlying state of the world $\omega \in \{1, \ldots k\}$ such that θ_ℓ and θ_r are independent given ω. That is, $\pi(\theta_\ell, \theta_r | \omega) = \pi_\ell(\theta_\ell | \omega)\pi_r(\theta_r | \omega) \quad \forall \theta_\ell, \theta_r, \omega$. Then

$$\pi(\theta_\ell, \theta_r) \; = \; \sum_{\omega=1}^{k} \pi_\ell(\theta_\ell|\omega)\pi_r(\theta_r|\omega)\,\Pr(\omega).$$

and so each column $\pi(\cdot, \theta_r)$ of the matrix $(\pi(\theta_\ell, \theta_r))_{\theta_\ell, \theta_r}$ is a linear combination of the k vectors $\pi_\ell(\cdot|\omega)$, $\omega = 1, \ldots, k$ (with coefficients $\alpha_{\theta_r}(\omega) = \pi_r(\theta_r|\omega)\Pr(\omega)$). Hence $\mathrm{rank}(\pi) \leq k$.

terms of the principal's objective and applies when the principal is ex-ante indifferent between the two options. The second characterization is in terms of a related optimal transport problem. It also yields an explicit characterization of incentive-compatible mechanisms under independence.

2.5.1 The role of the objective

Definition 2.2. The principal's objective is said to be *additive* if there exist functions $v_i : \Theta_i \to \mathbb{R}$ such that

$$v(\theta_\ell, \theta_r) \;=\; v_\ell(\theta_\ell) + v_r(\theta_r) \quad \forall \theta_\ell, \theta_r.$$

Given $\pi \in \Pi$ the objective is said to be *π-additive* if there exist coefficients $\lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)$, $\lambda_r(\theta_r, \tilde{\theta}_r) \in \mathbb{R}$ such that

$$v(\theta_\ell, \theta_r)\pi(\theta_\ell, \theta_r) \;=\; \sum_{\tilde{\theta}_\ell} \lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)\pi(\theta_r|\tilde{\theta}_\ell) + \sum_{\tilde{\theta}_r} \lambda_r(\theta_r, \tilde{\theta}_r)\pi(\theta_\ell|\tilde{\theta}_r) \quad \forall \theta_\ell, \theta_r. \tag{2.2}$$

Additivity is a special case of π-additivity (take $\lambda_i(\theta_i, \tilde{\theta}_i) = v_i(\theta_i)\pi_i(\tilde{\theta}_i)\mathbf{1}_{(\tilde{\theta}_i=\theta_i)}$) and it is easily seen that the two concepts coincide when types are independent ($\pi = \pi_\ell \pi_r$). To interpret π-additivity let the type distribution by $\pi \in \Pi$ and consider some mechanism x. When v is π-additive we then get from (2.2) that

$$E_\pi[v(\boldsymbol{\theta})x(\boldsymbol{\theta})] \;=\; \sum_{\theta_\ell, \tilde{\theta}_\ell} \lambda_\ell(\theta_\ell, \tilde{\theta}_\ell) E_\pi[x(\theta_\ell, \boldsymbol{\theta}_r)|\tilde{\theta}_\ell] + \sum_{\theta_r, \tilde{\theta}_r} \lambda_r(\theta_\ell, \tilde{\theta}_r) E_\pi[x(\boldsymbol{\theta}_\ell, \theta_r)|\tilde{\theta}_r]$$

so that $E_\pi[v(\boldsymbol{\theta})x(\boldsymbol{\theta})]$ is a linear combination of the potential expected payoffs $E_\pi[x(\theta_i, \boldsymbol{\theta}_{-i})|\tilde{\theta}_i]$ of types $\tilde{\theta}_i$ from any (mis-)report θ_i. If x is IC then $E_\pi[v(\boldsymbol{\theta})x(\boldsymbol{\theta})]$ is the principal's expected payoff from x and the "misreporting expectations" must all coincide with the maximin value $\bar{x}$. Hence replacing x by the constant mechanism $\tilde{x} \equiv \bar{x}$ does not change the principal's payoff and x cannot be profitable. A necessary condition for the existence of a profitable mechanism is thus that the principal's objective is not π-additive. If the principal is ex-ante indifferent between L and R then this condition is also sufficient

Proposition 2.5. *Let types be distributed according to $\pi \in \Pi$.*

A profitable mechanism exists only if the principal's objective is not π-additive.

If $E_\pi[v(\boldsymbol{\theta})] = 0$, then a profitable mechanism exists if and only if the principal's objective is not π-additive. In particular, if $E_\pi[v(\boldsymbol{\theta})] = 0$ and types are independent, then a profitable mechanism exists if and only if the principal's objective is not additive.

The proof (in the appendix) works by projecting $v\pi$ on the linear subspace U of functions that can be expressed in the form of the right hand side of (2.2). Given ex-ante indifference the principal's expected payoff in an IC mechanism depends only on the part of $v\pi$ that is orthogonal to U. We construct a mechanism that yields a strictly positive payoff whenever this projection residual is nonzero.[5]

Example 2.1 (continued). Consider again the budget allocation problem. Recall that

$$v(\theta_\ell, \theta_r) \;=\; E[y_L - y_R | \theta_\ell, \theta_r] \;=\; E[y_L | \theta_\ell, \theta_r] - E[y_R | \theta_\ell, \theta_r].$$

Let $h_i(\theta_\ell, \theta_r) = E[y_i | \theta_i, \theta_{-i}]$. If $h_i(\theta_\ell, \theta_r)$ depends only on θ_i then Proposition 2.5 implies that there does not exist a profitable mechanism. Hence a necessary condition for the existence of a profitable mechanism is that at least one department head has information that is relevant to the future return of the other department. Now assume that types are independent and identically distributed, and that $h_\ell = h_r = h$. Then $E[y_\ell] = E[h(\theta_\ell, \theta_r)] = E[h(\theta_r, \theta_\ell)] = E[y_r]$ so that the principal is ex-ante indifferent. Hence a profitable mechanism exists if, and only if $h(\theta_\ell, \theta_r) - h(\theta_r, \theta_\ell)$ is not additive.

2.5.2 The role of correlation

Correlation between agent-types affects the principal through two distinct channels. Firstly, correlation affects the set of mechanisms in which agents find it optimal to be truthful (see Section 2.4). Secondly, fixing a mechanism and assuming that agents are truthful, correlation can increase or decrease the principal's expected payoff by concentrating more mass on specific type profiles. In this section we show that the principal's problem can be viewed as a problem of choosing an "optimal correlation structure".

We start by reinterpreting incentive-compatibility. The proof is in the appendix.

Lemma 2.6. *Let the type distribution be $\pi \in \Pi$. A mechanism x is IC if and only if*

(1) Agents are ex-ante indifferent between reports; that is,

$$E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})] \;=\; E_\pi[x(\theta_i'', \boldsymbol{\theta}_{-i})] \quad \forall \theta_i', \theta_i'' \,\forall i.$$

(2) Their type realizations are uninformative; that is,

$$E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i}) | \theta_i] \;=\; E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})] \quad \forall \theta_i, \theta_i' \,\forall i.$$

Ex-ante indifference is equivalent to IC under the independent type distribution $\pi_\ell \pi_r$. Uninformativeness implies that agents cannot gain any payoff-relevant information from their type about their opponent's type. Note that this is automatically satisfied if types are independent. Correlation therefore restricts the set of IC

5. One appeal of the argument given in the appendix is that it is constructive. There is an alternative and more abstract argument using Farkas' lemma (though we do not report this argument here).

mechanisms by making the agents more informed which adds additional incentive-constraints. From this perspective, IC under correlation lies mid-way between IC under independence and IC under full information.

Lemma 2.6 allows us to derive a necessary and sufficient criterion for the existence of a profitable mechanism. We need the following definition.

Definition 2.3. Two joint type distributions $\pi, \tilde{\pi} \in \Pi(\pi_\ell, \pi_r)$ with the same marginals $\pi_i > 0$ are said to be orthogonal if

$$\mathrm{Cov}(\pi(\theta_i|\boldsymbol{\theta}_{-i}), \tilde{\pi}(\theta_i'|\boldsymbol{\theta}_{-i})) = 0 \quad \forall \theta_i, \theta_i' \, \forall i.$$

Hence π and $\tilde{\pi}$ are orthogonal if the random variables $\pi(\theta_i|\boldsymbol{\theta}_{-i})$ and $\tilde{\pi}(\theta_i'|\boldsymbol{\theta}_{-i})$ are uncorrelated for all θ_i, θ_i', $i = \ell, r$. Note that

$$\mathrm{Cov}(\pi(\theta_i|\boldsymbol{\theta}_{-i}), \tilde{\pi}(\theta_i'|\boldsymbol{\theta}_{-i})) = \sum_{\theta_{-i}} [\pi(\theta_i|\theta_{-i}) - \pi_i(\theta_i)][\tilde{\pi}(\theta_i'|\theta_{-i}) - \pi_i(\theta_i')]\pi_{-i}(\theta_{-i})$$

and $\pi(\theta_i|\theta_{-i}) - \pi_i(\theta_i)$ is the update of type θ_{-i} about the probability of type θ_i under π. Clearly, if one of π or $\tilde{\pi}$ is the independent type distribution $\pi_l \pi_r$ then orthogonality is automatically satisfied. Otherwise the condition says that for all θ_i, θ_i', the vector $\pi(\theta_i|\cdot) - \pi_i(\theta_i) \in \mathbb{R}^{\Theta_{-i}}$ of possible belief updates of agent $-i$ about the probability of type θ_i under π must be orthogonal to the vector of updates $\tilde{\pi}(\theta_i'|\cdot) - \pi_i(\theta_i) \in \mathbb{R}^{\Theta_{-i}}$ about the probability of θ_i' under $\tilde{\pi}$ under the inner product $\langle a, b \rangle = \sum_{\theta_{-i}} a(\theta_{-i}) b(\theta_{-i})$ on $\mathbb{R}^{\Theta_{-i}}$.

The next result shows that the problem of finding a profitable mechanism is intricately related to the choice of an "optimal correlation strucuture": A profitable mechanism exists if and only if it is possible to find some alternative correlation structure that is orthogonal to the exogenously given one and such that — under the alternative correlation structure (and with a suitably transformed objective) — L becomes the principal's ex-ante strictly preferred option. This can be phrased as a constrained optimal transport problem.[6]

Proposition 2.7. *Let the type distribution be $\pi \in \Pi$ and denote its marginals by π_i. Let $\hat{v} = v\pi/\pi_\ell \pi_r$. A profitable mechanism exists if and only if*

$$\left(\max_{\tilde{\pi} \in \Pi(\pi_\ell, \pi_r)} E_{\tilde{\pi}}[\hat{v}(\boldsymbol{\theta})] \text{ s.t. } \tilde{\pi} \text{ is orthogonal to } \pi \right) > 0. \tag{2.3}$$

In particular, if types are independent then a profitable mechanism exists if, and only if

$$\max_{\tilde{\pi} \in \Pi(\pi_\ell, \pi_r)} E_{\tilde{\pi}}[v(\boldsymbol{\theta})] > 0.$$

6. For an in-depth treatment of optimal transport see Villani (2009).

To explain how the constrained optimal transport problem in Proposition 2.7 is related to the principal's problem let x be some mechanism. Together with π, x induces a density $g(\theta) = \pi(\theta)x(\theta) = \frac{\pi(\theta)}{\pi_\ell(\theta_\ell \pi_r(\theta_r)}f(\theta)$ of a measure on Θ whose "correlation structure" depends on an exogenous part $\frac{\pi}{\pi_\ell \pi_r}$ and an endogenous part f. Instead of in terms of mechanisms, the principal's problem can also be phrased in terms of f. Requiring ex-ante indifference for the agents then translates into requiring that f be a nonnegative multiple of some probability distribution $\tilde{\pi} \in \Pi$ with the same marginals as π: $f = q\tilde{\pi}$ $(q \in [0, 1])$. Uninformativeness translates into $\tilde{\pi}$ being orthogonal to π. Under this reparametrization the principal's objective becomes $qE_{\tilde{\pi}}[\hat{v}(\theta)]$ and the (upper) feasibility constraint on the mechanism becomes a correlation constraint: $q\frac{\tilde{\pi}}{\pi_\ell \pi_r} \leq 1$. If there is a profitable $\tilde{\pi}$ then the principal therefore faces a tradeoff between up-scaling her objective (by increasing q) and the ability to concentrate more mass on type profiles with a positive objective value (by decreasing q). The mere existence of a profitable $\tilde{\pi}$ does not depend on the correlation constraint, however, and after dropping this constraint and dividing everything by q we arrive at the formulation in Proposition 2.7.

The proof of the above proposition (in the appendix) yields another characterization of incentive compatible mechanisms when types are independent.

Corollary 2.8. *If types are independent then a mechanism x is IC if and only if there exist nonnegative coefficients $\{\gamma_j\}_{j=1}^k$ $(k \geq 0)$ and extreme points*[7] *π^j of $\Pi(\pi_\ell, \pi_r)$ such that*

$$x = \sum_{j=1}^{k} \frac{\pi^j}{\pi_\ell \pi_r} \gamma_j.$$

Consider an example where both agents have the same number of types (without loss $\Theta_\ell = \Theta_r$) and where marginals are uniform. Together with the Birkhoff-von Neumann Theorem the characterization then implies that a mechanism is IC if and only if it can be decomposed into mechanisms where, up to relabeling of the types, the principal chooses L if and only if both agents make the same report. This illustrates how incentive-compatibility is fundamentally based on the inability (and unwillingness) of the agents to coordinate.

Example 2.5. Assume $\Theta_\ell = \Theta_r = \{1, \ldots, m\}$. A mechanism x is said to be a match-your opponent mechanism if there exists a matching[8] $m : \Theta_\ell \to \Theta_r$ such that

$$x(\theta_\ell, \theta_r) = \begin{cases} 1, & \text{if } \theta_r = m(\theta_\ell) \\ 0, & \text{otherwise} \end{cases}$$

7. Recall that an element of a convex set is an extreme point of the set if is is not the midpoint of a line-segment connecting two distinct points in the set. For a characterization of the extreme points of $\Pi(\pi_\ell, \pi_r)$ see Brualdi (2006), Theorem 8.1.2.

8. A matching is a bijective function.

Assume that types are independent with $\pi_i = \frac{1}{N}$. Using the Birkhoff-von Neumann Theorem and Corollary 2.8, a mechanism x is IC if and only if there exist match-your-opponent mechanisms x^j and nonnegative coefficients γ_j such that

$$x = \sum_{j=1}^{k} x^j \gamma_j.$$

Thus, a profitable mechanism exists if and only if there exists a profitable match-your-opponent mechanism. If the principal's objective is supermodular it follows[9] that a profitable mechanism exists if and only if

$$\sum_{t=1}^{m} v(t,t) > 0.$$

Indeed, as long as types are independent and agents are symmetric (i.e. $\pi_\ell(t) = \pi_r(t)$, $t = 1, \ldots, m$), it can be shown that a profitable mechanism exists if and only if $\sum_{t=1}^{m} \pi_\ell(t) v(t,t) > 0$.[10]

2.6 Allocation with more than two agents and disposal

One application of our results is the problem of allocating a single nondisposable good between two agents. In this section, we extend our setting and study the problem of allocating a (potentially disposable) good among n agents $i = 1, \ldots, n$.

Agents are again expected utility maximizers and enjoy utility 1 from receiving the good and 0 otherwise.[11] Every agent has a private type $\theta_i \in \Theta_i$. The set of type profiles $\Theta = \Pi_i \Theta_i$ is finite. Throughout this section we assume that types are independent; the joint type distribution is denoted by $\pi(\theta_1, \ldots, \theta_n) = \pi_1(\theta_1) \ldots \pi_n(\theta_n)$[12].

The principal's value from allocating the good to agent i can depend on the types $\theta = (\theta_1, \ldots, \theta_n)$ of all agents and is denoted by $v_i(\theta) \in \mathbb{R}$. A (direct) mechanism specifies for each agent i and every profile θ the probability of allocating the good to this agent when the report profile is θ.

We distinguish between the case where the principal can commit to dispose the good (or consume it privately) from the case where she is forced to allocate to one of the agents. We normalize the principal's utility from disposing the good to 0. If the principal must allocate the good the feasibility constraint reads: $\sum_{i=1}^{n} x_i(\theta) = 1$; under free disposal it reads: $\sum_{i=1}^{n} x_i(\theta) \leq 1$. In either case, the principal's problem is to

9. See Hardy, Littlewood, and Pólya (1952), Becker (1973), and Vince (1990).

10. See Hoffman (1963).

11. All results apply unchanged if agents receive utility $\bar{u}_i(\theta_i) > 0$ from getting the good and 0 otherwise.

12. As before, we assume without loss of generality that $\pi_i > 0$, $i = 1, \ldots, n$.

find a feasible, incentive compatible mechanism that maximizes $E[\sum_{i=1}^{n} v_i(\boldsymbol{\theta})x_i(\boldsymbol{\theta})]$. As before, we will be interested in whether the principal can do better than choosing her ex-ante preferred option.

We first characterize the set of incentive compatible mechanisms. Whether or not disposal is possible, a mechanism is incentive compatible if and only if each agent's interim probability of obtaining the good does not depend on his report:

Lemma 2.9. *Assume there are n agents with independent types and let x be a mechanism (with or without disposal). Then x is incentive compatible if and only if*

$$E[x_i(\theta_i, \boldsymbol{\theta}_{-i})] = E[x_i(\boldsymbol{\theta})] \quad \forall i \, \forall \theta_i.$$

Let $\bar{v} = \max_i E[v_i(\boldsymbol{\theta})]$ be the principal's expected payoff from allocating to her ex-ante preferred agent.

An incentive compatible mechanism is profitable if it yields the principal strictly more than choosing her ex-ante preferred option (ignoring type reports). Formally, when there is free disposal, an incentive compatible mechanism is profitable if $\sum_i E[v_i(\boldsymbol{\theta})x_i(\boldsymbol{\theta})] > \max\{0, \bar{v}\}$. Without disposal it is profitable if $\sum_i E[v_i(\boldsymbol{\theta})x_i(\boldsymbol{\theta})] > \bar{v}$.

The following proposition generalizes the scope of Proposition 2.5 to the n agent case under independence.

Proposition 2.10. *Assume there are n agents with independent types and that the principal has to allocate to some agent. If the principal is unbiased then a profitable mechanism exists if and only if there do not exist functions $u_1(\theta_1), \ldots, u_n(\theta_n)$ such that*[13]

$$v_i(\theta) - v_j(\theta) = u_i(\theta_i) - u_j(\theta_j) \quad \forall i, j \, \forall \theta. \tag{2.4}$$

In the proof (in the appendix) we show that a violation of condition (2.4) remains a necessary condition for the existence of a profitable mechanism when the principal is not unbiased. The above proposition also allows us to state a necessary and sufficient condition for the existence of a profitable mechanism when the principal is allowed to discard the good:

Proposition 2.11. *Assume there are n agents with independent types and that the principal does not have to allocate the good to the agents. If the principal is unbiased*

13. Equivalently: There do not exist functions $\omega(\theta)$ and $u_1(\theta_1), \ldots, u_n(\theta_n)$ such that

$$v_i(\theta) = \omega(\theta) + u_i(\theta_i) \quad \forall i, \theta.$$

and $\bar{v} = 0$, then a profitable mechanism exists if and only if there do not exist functions $u_1, \ldots, u_n$ such that

$$v_i(\theta) = u_i(\theta_i) \quad \forall i \forall \theta;$$

that is, a profitable mechanisms exists if and only if there is an agent j and a type $\theta_j \in \Theta_j$ such that $v_j(\theta_j, \theta_{-j})$ is non-constant in θ_{-j}.

2.7 Related Literature

Our main setting can be interpreted as an allocation problem without disposal. It therefore relates to Myerson (1981) who characterizes the set of IC-mechanisms with transfers under independence. Crémer and McLean (1985, 1988) show that with transfers, full rank correlation makes any allocation rule implementable.

Börgers and Postl (2009) study a setting without transfers and two agents with opposed interests. Their setting has a third option that acts as a compromise and types are iid. They consider utilitarian welfare and study second-best rules using numerical tools. Since utilitarian welfare is additive, our results underline the importance of the compromise option for their results. Kim (2017) considers a related setting with at least three ex-ante symmetric alternatives and several agents with iid private values whose interests are not necessarily opposed.

Feng and Wu (2019) ask in a setting without transfer with a perfect conflict of interests not between the agents but between the agents and the principal if the later can do better than choosing her ex-ante preferred option. Goldlücke and Tröger (2020) study "threshold mechanisms" with binary message spaces to assign an unpleasant task without transfers in a setting with symmetric agents with iid types.

More broadly, our paper is related to a strand of the mechanism design literature investigating allocation problems without transfers. Several papers in this literature study non-monetary instruments for screening the agents when the agents' types are independent and the principal's objective is additive (see, for example, Ben-Porath, Dekel, and Lipman, 2014; Mylovanov and Zapechelnyuk, 2017). Kattwinkel and Knoepfle (2019) and Kattwinkel (2020) consider a single-agent allocation problems where the principal observes a private signal about the agent's type. Niemeyer and Preusser (2022) consider *dominant-strategy* IC mechanisms with correlated types.

The proof of Proposition 2.2 connects implementation of mechanisms with the properties of correlated equilibrium Aumann (1974, 1987) in zero sum games Rosenthal (1974).

Our comparative statics result for the set of IC mechanisms with respect to the spanning preorder relates to the comparison of experiments (Blackwell, 1951, 1953; see also Börgers, Hernando-Veciana, and Krähmer, 2013), and the comparison of information structures in games (Bergemann and Morris, 2016).

Since the set of DIC mechanisms in our setting coincides with the set of constant mechanisms, the existence of non-constant IC mechanisms is related to BIC-DIC equivalence Manelli and Vincent (2010) and Gershkov et al. (2013). In our setting, IC under correlation can be viewed as a mid-way point between IC under independence and DIC.

Appendix 2.A Omitted proofs

2.A.1 Comparative statics for implementation

2.A.1.1 Proof of Proposition 2.4

Proof. The first assertion is an immediate consequence of Proposition 2.3 and Example 2.2. To see why the second assertion holds note that if $\tau^0, \tau^1, \dots, \tau^k \in \Delta\Theta_{-i}$ and $y(\theta_{-i}) : \Theta_{-i} \to \mathbb{R}$ is any function such that

$$\sum_{\theta_{-i}} \tau^j(\theta_{-i}) y(\theta_{-i}) \;=\; 0 \quad \forall j \;=\; 1, \dots, k.$$

then also $\sum_{\theta_{-i}} \tau^0(\theta_{-i}) y(\theta_{-i}) = 0$ if τ^0 is spanned by $\{\tau^1, \dots, \tau^k\}$.

Now assume without loss that $|\Theta_\ell| \geq |\Theta_r|$. By the full rank-condition the vectors $(\pi(\cdot|\theta_\ell))_{\theta_\ell}$ contain a basis of $\mathbb{R}^{\Theta_r}$. In particular, for any θ_r, they span the belief $\mathbf{1}_{\theta_r}$ which puts mass 1 on θ_r. But then ℓ's IC constraints must be satisfied under that belief (consider the function $y_{\theta_\ell'}(\theta_r) = x(\theta_\ell', \theta_r) - \bar{x}$). But that means that

$$x(\theta_\ell', \theta_r) \;=\; \bar{x} \quad \forall \theta_\ell'.$$

Since θ_r was arbitrary it follows that x must be constant.

Next we will show that a distribution $\pi \in \Pi$ is maximal iff it has full rank, and that π is minimal iff it is an independent type distribution. We need to show that (i) π has full rank iff it spans every $\tilde{\pi}$ that spans it and (ii) π is an independent type distribution iff for every $\tilde{\pi}$ spanned by π it is also the case that $\tilde{\pi}$ spans π. Throughout, assume without loss of generality that $|\Theta_\ell| \geq |\Theta_r|$.

First note that for $\pi, \tilde{\pi} \in \Pi$, π spans $\tilde{\pi}$ if and only if the row space and the column space of $\tilde{\pi}$ are contained in the row space and the column space, respectively, of π.

Assume that $\pi \in \Pi$ has full rank and that $\tilde{\pi} \in \Pi$ spans π. Denote the column and row spaces of π and $\tilde{\pi}$ by V, W and $\tilde{V}, \tilde{W}$, respectively. Since $\tilde{\pi}$ spans π it holds that $V \subset \tilde{V}$ and $W \subset \tilde{W}$. Since π has full rank, V and W both have dimension $|\Theta_r|$. $\tilde{\pi}$ is an $|\Theta_\ell| \times |\Theta_r|$ matrix and so its column and row spaces cannot have a dimension larger than $|\Theta_r|$. Hence we must have $V = \tilde{V}$ and $W = \tilde{W}$. But that implies that π also spans $\tilde{\pi}$. Thus π is maximal. Conversely, assume that π is maximal. Let $\tilde{\pi} \in \Pi$ be some full rank distribution that spans π. Since π is maximal, π must then also span $\tilde{\pi}$. Hence the row space (and also the column space) of π must have dimension $|\Theta_r|$. This means that π has full rank.

Now let π be an independent distribution. Let $\tilde{\pi}$ be spanned by π. Then, for all θ_i, $\tilde{\pi}(\cdot|\theta_i)$ is a linear combination of the vectors $\pi(\cdot|\tilde{\theta}_i)$ ($\tilde{\theta}_i \in \Theta_i$). But since types are independent under π, the latter vectors all coincide with $\pi_i(\cdot)$. Hence $\tilde{\pi}(\cdot|\theta_i) = \pi_i(\cdot)$ for all θ_i, $i = \ell, r$. Thus $\tilde{\pi} = \pi$; in particular $\tilde{\pi}$ spans π. Hence π is a minimal element. Now let $\tilde{\pi}$ be a minimal element. Let $\tilde{\pi}$ be the independent distribution with the same marginals as π. Then π spans $\tilde{\pi}$ and since π is minimal, $\tilde{\pi}$ must also span π. But $\tilde{\pi}$ has rank one and so π must also have rank one. But that means that π must be an independent type distribution (thus $\pi = \tilde{\pi}$). □

2.A.2 Profitable mechanisms

2.A.2.1 Proof of Proposition 2.5

Proof. First consider a general $\pi \in \Pi$. Define $w(\theta) = v(\theta)\pi(\theta)$. If $E_\pi[v(\boldsymbol{\theta})] = 0$ and agents are truthful then the principal's payoff from some mechanism x is

$$\begin{aligned}\sum_\theta v(\theta)\pi(\theta)x(\theta) &= E_\pi[v(\boldsymbol{\theta})]E_\pi[x(\boldsymbol{\theta})] + \sum_\theta v(\theta)\pi(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]) \\ &= \sum_\theta w(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]).\end{aligned}$$

By Proposition 2.2, x is IC if and only if

$$\sum_{\theta_{-i}} \pi(\theta_{-i}|\theta_i)(x(\theta_i', \theta_{-i}) - E_\pi[x(\boldsymbol{\theta})]) = 0 \quad \forall \theta_i, \theta_i' \, \forall i. \tag{3}$$

First assume that there exist coefficients $\lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)$, $\lambda_r(\theta_r, \tilde{\theta}_r) \in \mathbb{R}$ such that

$$w(\theta) = \sum_{\tilde{\theta}_\ell} \lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)\pi(\theta_r|\tilde{\theta}_\ell) + \sum_{\tilde{\theta}_r} \lambda_r(\theta_r, \tilde{\theta}_r)\pi(\theta_\ell|\tilde{\theta}_r) \quad \forall \theta_\ell, \theta_r. \tag{4}$$

Then by (3), any IC mechanism satisfies $\sum_\theta w(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]) = 0$ and so there is no profitable mechanism.

Now assume instead that there exist no coefficients λ_ℓ and λ_r such that w satisfies (4). Let U be the set of all $u \in \mathbb{R}^{\Theta_\ell \times \Theta_r}$ for which there exist coefficients $\lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)$, $\lambda_r(\theta_r, \tilde{\theta}_r) \in \mathbb{R}$ such that

$$u(\theta) = \sum_{\tilde{\theta}_\ell} \lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)\pi(\theta_r|\tilde{\theta}_\ell) + \sum_{\tilde{\theta}_r} \lambda_r(\theta_r, \tilde{\theta}_r)\pi(\theta_\ell|\tilde{\theta}_r) \quad \forall \theta_\ell, \theta_r.$$

Note that U is a linear subspace of $\mathbb{R}^{\Theta_\ell \times \Theta_r}$ (Indeed, $0 \in U$ and U is closed under addition and multiplication by scalars). Let $\hat{u}$ be the orthogonal projection of w onto U and let $\hat{w}$ be the orthogonal projection of w onto the orthogonal complement of U so that $w = \hat{w} + \hat{u}$. By assumption $w \notin U$, and so $\hat{w} \neq 0$.

As before, given any IC mechanism x it holds that $\sum_\theta \hat{u}(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]) = 0$ and so the principal's payoff from any IC mechanism x is

$$\sum_\theta w(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]) \;=\; \sum_\theta \hat{w}(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]).$$

We will now construct a profitable mechanism, i.e. an IC mechanism for which the latter expression is positive. Define

$$\hat{x}(\theta) \;=\; \varepsilon(\hat{w}(\theta) - \min \hat{w})$$

where $\varepsilon > 0$ is sufficiently small such that $\hat{x} \leq 1$. First note that $\hat{x}$ is IC. Indeed, for any $\theta'_\ell, \theta''_\ell$,

$$\begin{aligned}
\sum_{\theta_r} \pi(\theta_r|\theta''_\ell)\hat{x}(\theta'_\ell, \theta_r) \;&=\; \varepsilon \sum_{\theta_r} \pi(\theta_r|\theta''_\ell)\hat{w}(\theta'_\ell, \theta_r) - \varepsilon \min \hat{w} \\
&=\; \varepsilon \sum_{\theta_\ell}\sum_{\theta_r} \underbrace{\left(\sum_{\tilde{\theta}_\ell} \mathbf{1}_{\theta_\ell = \theta'_\ell, \tilde{\theta}_\ell = \theta''_\ell} \pi(\theta_r|\tilde{\theta}_\ell)\right)}_{\in U} \hat{w}(\theta_\ell, \theta_r) - \varepsilon \min \hat{w} \\
&=\; -\varepsilon \min \hat{w},
\end{aligned}$$

where the last equality follows because $\hat{w}$ lies in the orthogonal complement of U. Hence $\hat{x}$ is IC for agent ℓ. The proof that $\hat{x}$ is IC for agent r is symmetric.

We now show that $\hat{x}$ is actually a profitable mechanism. First note that the above incentive-compatibility calculation implies that $E_\pi[\hat{x}(\boldsymbol{\theta})|\theta_\ell] = -\varepsilon \min \hat{w}$ and in particular $E_\pi[x(\boldsymbol{\theta})] = -\varepsilon \min \hat{w}$. Thus the principal's payoff from $\hat{x}$ is

$$\begin{aligned}
\sum_\theta \hat{w}(\theta)(x(\theta) - E_\pi[x(\boldsymbol{\theta})]) \;&=\; \varepsilon \sum_\theta \hat{w}(\theta)\hat{w}(\theta) \\
&>\; 0.
\end{aligned}$$

Hence $\hat{x}$ is a profitable mechanism.

Finally, assume that types are independent. Note that

$$v(\theta_\ell, \theta_r)\pi_\ell(\theta_\ell)\pi_r(\theta_r) \;=\; \sum_{\tilde{\theta}_\ell} \lambda_\ell(\theta_\ell, \tilde{\theta}_\ell)\pi(\theta_r) + \sum_{\tilde{\theta}_r} \lambda_r(\theta_r, \tilde{\theta}_r)\pi(\theta_\ell) \quad \forall \theta_\ell, \theta_r$$

if and only if

$$v(\theta_\ell, \theta_r) \;=\; \underbrace{\sum_{\tilde{\theta}_\ell} \tilde{\lambda}_\ell(\theta_\ell, \tilde{\theta}_\ell)}_{v_\ell(\theta_\ell)} + \underbrace{\sum_{\tilde{\theta}_r} \tilde{\lambda}_r(\theta_r, \tilde{\theta}_r)}_{v_r(\theta_r)},$$

where $\tilde{\lambda}_i(\theta_i, \tilde{\theta}_i) = \frac{\lambda_i(\theta_i, \tilde{\theta}_i)}{\pi_i(\theta_i)}$ and so the earlier condition reduces to additivity. $\square$

2.A.2.2 Proof of Lemma 2.6

Proof. Let $\pi \in \Pi$ and let x be IC. By Proposition 2.2, agents must be ex-ante indifferent between reports and their type realizations must be uninformative. Conversely, suppose that x satisfies the assumptions of Lemma 2.6. Ex-ante indifference combined with the law of iterated expectations implies that $E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})] = E_\pi[x(\boldsymbol{\theta})]\ \forall i, \theta_i'$. Hence for any i and θ_i, θ_i':

$$\begin{aligned} E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})|\theta_i] &= E_\pi[x(\theta_i', \boldsymbol{\theta}_{-i})] \\ &= E_\pi[x(\boldsymbol{\theta})], \end{aligned}$$

where the first equality follows from uninformativeness. □

2.A.2.3 Proof of Proposition 2.7

Proof. By Lemma 2.6, the principal's problem can be written as

$$\begin{aligned} \max \quad & \sum_\theta \hat{v}(\theta)\pi_\ell(\theta_\ell)\pi_r(\theta_r)x(\theta) \\ \text{s.t.} \quad & \sum_{\theta_{-i}} \pi_{-i}(\theta_{-i})x(\theta_i', \theta_{-i}) = \sum_\theta \pi_\ell(\theta_\ell)\pi_r(\theta_r)x(\theta_\ell, \theta_r) && \forall \theta_i'\ \forall i && (I_i) \\ & \sum_{\theta_{-i}} \pi(\theta_{-i}|\theta_i)x(\theta_i', \theta_{-i}) = \sum_{\theta_{-i}} \pi_{-i}(\theta_{-i})x(\theta_i', \theta_{-i}) && \forall \theta_i, \theta_i'\ \forall i && (U_i) \\ & 0 \leq x(\theta) \leq 1 && \forall \theta. && (F) \end{aligned}$$

Here, (I_i) are the ex-ante indifference constraints (or, equivalently, the IC constraints under the independent type distribution $\pi_\ell \pi_r$) and (U_i) are the uninformativeness constraints. Now define

$$f(\theta_\ell, \theta_r) = \pi_\ell(\theta_\ell)\pi_r(\theta_r)x(\theta_\ell, \theta_r).$$

Using this substitution the principal's objective becomes $\sum_\theta v(\theta)f(\theta)$, the ex-ante indifference constraints become

$$\sum_{\theta_{-i}} f(\theta_i', \theta_r) = \pi_i(\theta_i')\sum_\theta f(\theta) \qquad \forall \theta_i'\ \forall i \tag{2.A.2}$$

and the uninformativeness constraints can be written as

$$\sum_{\theta_{-i}} (\pi(\theta_i|\theta_{-i}) - \pi_i(\theta_i))f(\theta_i', \theta_{-i}) = 0 \qquad \forall \theta_i, \theta_i'\ \forall i \tag{2.A.3}$$

while the feasibility constraints become

$$0 \leq f(\theta_\ell, \theta_r) \leq \pi_\ell(\theta_\ell)\pi_r(\theta_r) \qquad \forall \theta_\ell, \theta_r.$$

Equation (2.A.2) says that the marginals of f are proportional to π_i. Since f is nonzero, it is thus a nonnegative multiple of some joint probability distribution $\tilde{\pi}$ with marginals π_i. Hence the principal's problem can be written as

$$\begin{aligned}
\max_{q\in[0,1]} \max_{\tilde{\pi}\in\Pi(\pi_l,\pi_r)} \quad & q\sum_{\theta}\hat{v}(\theta)\tilde{\pi}(\theta) & \\
\text{s.t.} \quad & \sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))\tilde{\pi}(\theta_i',\theta_{-i}) = 0 & \forall\theta_i,\theta_i'\,\forall i \\
& q\tilde{\pi}(\theta_\ell,\theta_r) \le \pi_l(\theta_\ell)\pi_r(\theta_r) & \forall\theta_\ell,\theta_r
\end{aligned}$$

where $\Pi(\pi_l,\pi_r)$ is the set of joint type distributions with marginals π_i. A profitable mechanism therefore exists if and only if the latter problem has a positive optimal value.

Since any $\tilde{\pi}\in\Pi(\pi_\ell,\pi_r)$ can be made to satisfy the constraint $q\tilde{\pi}(\theta_\ell,\theta_r)\le\pi_l(\theta_\ell)\pi_r(\theta_r)$ after appropriate scaling, the problem's optimal value is positive if and only if the value of the relaxed problem in which that constraint is left out is positive. Finally, note that

$$\begin{aligned}
\sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))\tilde{\pi}(\theta_i',\theta_{-i}) &= \sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))\tilde{\pi}(\theta_i'|\theta_{-i})\pi_{-i}(\theta_{-i}) \\
&= \sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))(\tilde{\pi}(\theta_i'|\theta_{-i})-\pi_i(\theta_i'))\pi_{-i}(\theta_{-i})
\end{aligned}$$

because $\sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))\pi_i(\theta_i')\pi_{-i}(\theta_{-i})=0$. Therefore a profitable $\tilde{\pi}$ exists if and only if the optimal value of the following problem is positive:

$$\begin{aligned}
\max_{\tilde{\pi}\in\Pi(\pi_l,\pi_r)} \quad & \sum_{\theta}\hat{v}(\theta)\tilde{\pi}(\theta) & \\
\text{s.t.} \quad & \sum_{\theta_{-i}}(\pi(\theta_i|\theta_{-i})-\pi_i(\theta_i))(\tilde{\pi}(\theta_i'|\theta_{-i})-\pi_i(\theta_i'))\pi_{-i}(\theta_{-i}) = 0 & \forall\theta_i,\theta_i'\,\forall i.
\end{aligned}$$

This concludes the proof of Proposition 2.7. □

2.A.2.4 Proof of Corollary 2.8

Proof. The proof of Proposition 2.7 shows that under independence, a mechanism x is IC if and only if there exists some $q\in[0,1]$ and $\tilde{\pi}\in\Pi(\pi_l,\pi_r)$ such that $\pi_l\pi_r x = q\tilde{\pi}$. The set $\Pi(\pi_l,\pi_r)$ is a polytope (known as the transportation polytope), hence by the Weyl-Minkowski Theorem it is the convex hull of its finitely many extreme points. This implies the claim. □

2.A.3 Allocation with more than two agents and disposal

2.A.3.1 Proof of Lemma 2.9

Let $x = (x_1, \dots, x_n)$ be an incentive compatible mechanism. Let i be an agent and let $\theta_i, \theta_i' \in \Theta_i$. In order for type θ_i to be truthful, it must hold that $E[x_i(\theta_i, \boldsymbol{\theta}_{-i})] \geq E[x_i(\theta_i', \boldsymbol{\theta}_{-i})]$. In order for type θ_i' to be truthful, $E[x_i(\theta_i, \boldsymbol{\theta}_{-i})] \leq E[x_i(\theta_i', \boldsymbol{\theta}_{-i})]$ must hold. Hence in any incentive compatible mechanism $E[x_i(\theta_i, \boldsymbol{\theta}_{-i})]$ is constant in θ_i, for all i. Since any mechanism satisfying the latter is also incentive compatible, the condition is equivalent to incentive compatibility. Finally, if $E[x_i(\theta_i, \boldsymbol{\theta}_{-i})]$ is constant in θ_i then it must equal $E[x_i(\boldsymbol{\theta})]$.

2.A.3.2 Proof of Proposition 2.10

Proof. First assume that (2.4) holds. It follows that there exist functions $u_i(\theta_i)$ such that $v_i(\theta) - v_n(\theta) = u_i(\theta_i) - u_n(\theta_n)$ for all i and θ. Recall that in an IC mechanism x, $\bar{x}_i := E[x_i(\theta_i, \boldsymbol{\theta}_{-i})]$ does not depend on θ_i. The principal's payoff from an incentive compatible mechanism x is therefore

$$
\begin{aligned}
\sum_i E[v_i(\boldsymbol{\theta})x_i(\boldsymbol{\theta})] &= \sum_{i<n} E[(v_i(\boldsymbol{\theta}) - v_n(\boldsymbol{\theta}))x_i(\boldsymbol{\theta})] + E[v_n(\boldsymbol{\theta})\sum_i x_i(\boldsymbol{\theta})] \\
&= \sum_{i<n} E[(u_i(\boldsymbol{\theta}_i) - u_n(\boldsymbol{\theta}_n))x_i(\boldsymbol{\theta})] + E[v_n(\boldsymbol{\theta})] \\
&= \sum_{i<n} E[u_i(\boldsymbol{\theta}_i)E[x_i(\boldsymbol{\theta})|\boldsymbol{\theta}_i]] - E[u_n(\boldsymbol{\theta}_n)E[\sum_{i<n} x_i(\boldsymbol{\theta})|\boldsymbol{\theta}_n]] + E[v_n(\boldsymbol{\theta})] \\
&= \sum_{i<n} E[u_i(\boldsymbol{\theta}_i)]\bar{x}_i - E[u_n(\boldsymbol{\theta}_n)(1-\bar{x}_n)] + E[v_n(\boldsymbol{\theta})] \\
&= \sum_i E[u_i(\boldsymbol{\theta}_i)]\bar{x}_i + E[v_n(\boldsymbol{\theta}) - u_n(\boldsymbol{\theta}_n)] \\
&= \sum_i E[v_i(\boldsymbol{\theta}) + u_n(\boldsymbol{\theta}_n) - v_n(\boldsymbol{\theta})]\bar{x}_i + E[v_n(\boldsymbol{\theta}) - u_n(\boldsymbol{\theta}_n)] \\
&= \sum_i E[v_i(\boldsymbol{\theta})]\bar{x}_i.
\end{aligned}
$$

Hence, if (2.4) holds then the principal's expected payoff from an incentive compatible mechanism x is the same as her expected payoff from the constant mechanism y given by $y_i(\theta) \equiv \bar{x}_i$. In particular, the principal cannot do better than allocating to her ex-ante preferred agent and so no profitable mechanism exists. Note that we have not used the unbiasedness assumption and so the following is true even if the principal is not unbiased: A profitable mechanism can only exist if (2.4) is violated.

Now let the principal be unbiased. Assume that (2.4) is violated. Then there do no exist functions $u_i(\theta_i)$ $(i = 1, \dots, n)$ such that $v_i(\theta) - v_n(\theta) = u_i(\theta_i) - u_n(\theta_j)$ $(i < n)$. (If such functions did exist then it would follow that for any i,j: $v_i(\theta) - v_j(\theta) = v_i(\theta) - v_n(\theta) - (v_j(\theta) - v_n(\theta)) = u_i(\theta_i) - u_n(\theta_n) -$

$(u_j(\theta_j) - u_n(\theta_n)) = u_i(\theta_i) - u_j(\theta_j)$ and so (2.4) would hold). We will now construct a profitable mechanism.

Let Ω be the vector space of functions from $\{1, \ldots, n-1\} \times \Theta$ to $\mathbb{R}$ and let U_i be the set of functions from Θ_i to $\mathbb{R}$. Moreover, let $W \subset \Omega$ be the set of functions from $\{1, \ldots, n-1\} \times \Theta$ to $\mathbb{R}$ for which there exist functions $u_i(\theta_i)$ with $w_i(\theta) = \pi(\theta)(u_i(\theta_i) - u_n(\theta_n))\ \forall i < n\ \forall \theta$. Now consider the following minimization problem

$$\begin{aligned} &\min_{u_1 \in U_1, \ldots, u_n \in U_n} \sum_{i<n} \sum_{\theta} [\pi(\theta)(v_i(\theta) - v_n(\theta)) - \pi(\theta)(u_i(\theta) - u_n(\theta))]^2 \\ &= \min_{w \in W} \sum_{i<n} \sum_{\theta} [\tilde{v}_i(\theta) - w_i(\theta)]^2, \end{aligned}$$

where $\tilde{v}_i(\theta) = \pi(\theta)(v_i(\theta) - v_n(\theta))$. Note that W is a linear subspace of Ω and hence the solution $\hat{w}$ to the above minimization problem is the orthogonal projection of $\tilde{v} \in \Omega$ onto W (all spaces are finite-dimensional and so existence is not an issue). Let $\hat{\varepsilon} = \tilde{v} - \hat{w}$ be the projection residual. Note that the optimal value of the minimization problem is zero if and only if (2.4) holds. By assumption, (2.4) is violated and hence in particular $\hat{\varepsilon}$ must be nonzero. Moreover, since $\hat{\varepsilon}$ is orthogonal to W, for any $h \in W$ it must hold that

$$\sum_{i<n} \sum_{\theta} h_i(\theta) \hat{\varepsilon}(\theta) \; = \; 0.$$

We will now use $\hat{\varepsilon}$ to construct a profitable mechanism. Let $\underline{\hat{\varepsilon}} = \min_{i<n,\theta} \hat{\varepsilon}_i(\theta)$ and let

$$\hat{z}_i(\theta) \; = \; \hat{\varepsilon}_i(\theta) - \underline{\hat{\varepsilon}} \quad \forall i \; < \; n\, \forall \theta.$$

By construction, $\hat{z} \in \Omega$ is nonnegative. Define

$$\hat{x}_i(\theta) \; = \; \alpha \hat{z}_i(\theta),$$

where $\alpha > 0$ is chosen sufficiently small such that $\sum_{i<n} \hat{x}_i(\theta) \leq 1$ for all θ. Also, define $\hat{x}_n(\theta) = 1 - \sum_{i<n} \hat{x}_i(\theta)$. Then $\hat{x}$ is a feasible mechanism.

In the remainder of the proof we show that $\hat{x}$ is a profitable mechanism. We first verify that $\hat{x}$ is IC. Let $j < n$ be an agent. Then for any report θ_j' it holds that

$$\begin{aligned} \sum_{\theta_{-j}} \pi_{-j}(\theta_{-j}) \hat{x}_j(\theta_j', \theta_{-j}) \; &= \sum_{i<n} \sum_{\theta} \pi(\theta) \frac{1}{\pi_j(\theta_j)} 1(\theta_j \; = \; \theta_j') 1(i \; = \; j) \hat{\varepsilon}_i(\theta) - \alpha \underline{\hat{\varepsilon}} \\ &= \; -\alpha \underline{\hat{\varepsilon}}, \end{aligned}$$

because the function $\pi(\theta) \frac{1}{\pi_j(\theta_j)} 1(\theta_j = \theta_j') 1(i = j)$ lies in W and the function $\hat{\varepsilon}_i(\theta)$ is orthogonal to W. Since $-\alpha \underline{\hat{\varepsilon}}$ does not depend on θ_j', $\hat{x}$ is IC for agent j. It remains to check IC for agent n. Let θ_n' be a report. Then

$$\begin{aligned} \sum_{\theta_{-n}} \pi_{-n}(\theta_{-n}) \hat{x}_n(\theta_n', \theta_{-n}) \; &= \; \sum_{\theta_{-n}} \pi_{-n}(\theta_{-n})(1 - \sum_{i<n} \hat{x}_i(\theta_n', \theta_{-n})) \\ &= \; 1 + (n-1)\alpha \underline{\hat{\varepsilon}} - \alpha \sum_{i<n} \sum_{\theta} \pi(\theta) \frac{1}{\pi_n(\theta_n)} 1(\theta_n \; = \; \theta_n') \hat{\varepsilon}_i(\theta) \\ &= \; 1 + (n-1)\alpha \underline{\hat{\varepsilon}}, \end{aligned}$$

because the function $\pi(\theta)\frac{1}{\pi_n(\theta_n)}1(\theta_n = \theta'_n)$ lies in W and the function $\hat{\varepsilon}_i(\theta)$ is orthogonal to W. Hence $\hat{x}$ is an IC mechanism. It only remains to show that the principal's expected payoff from $\hat{x}$ is greater than $\hat{v}$.

The principal's expected payoff from $\hat{x}$ is

$$\begin{aligned}\sum_i \sum_\theta \pi(\theta) v_i(\theta) x_i(\theta) &= \sum_{i<n} \sum_\theta \pi(\theta)(v_i(\theta) - v_n(\theta))\hat{x}_i(\theta) + \sum_\theta \pi(\theta) v_n(\theta) \sum_i \hat{x}_i(\theta) \\ &= \sum_{i<n} \sum_\theta \tilde{v}_i(\theta)\hat{x}_i(\theta) + \bar{v} \\ &= \alpha \sum_{i<n} \sum_\theta (\hat{w}_i(\theta) + \hat{\varepsilon}_i(\theta))\hat{\varepsilon}_i(\theta) - \alpha \sum_i \sum_\theta \tilde{v}_i(\theta)\underline{\hat{\varepsilon}} + \bar{v}.\end{aligned}$$

By assumption, $\sum_\theta \pi(\theta) v_i(\theta)$ is the same for all i and hence for any $i < n$: $\sum_\theta \tilde{v}_i(\theta) = 0$. This means that the second term in the last line above is zero. Because in addition $\hat{w} \in W$ and $\hat{\varepsilon}$ is orthogonal to W, the principal's expected payoff now simplifies to

$$\alpha \sum_{i<n} \sum_\theta \hat{\varepsilon}_i(\theta)^2 + \bar{v}.$$

By assumption $\tilde{v}$ does not lie in W and so the projection residual $\hat{\varepsilon}$ is nonzero. It follows that the first term above is positive and therefore that the principal's expected payoff from $\hat{x}$ is greater than $\bar{v}$. That is to say, $\hat{x}$ is a profitable mechanism. □

2.A.3.3 Proof of Proposition 2.11

Proof. The result follows from Proposition 2.10 by interpreting the disposal option as an additional agent. Formally, let there be an agent 0 with a singleton type space $\Theta_0 = \{\theta^0\}$ and $v_0 \equiv 0$. A mechanism without disposal in this setting corresponds to a mechanism with disposal in the original setting. By Proposition 2.10, a profitable mechanism without disposal exists in the setting with the additional agent if and only if there do not exist functions $u_i(\theta_i)$ $(i = 0, \dots, n)$ such that $v_i(\theta) - v_0(\theta) = u_i(\theta_i) - u_0(\theta_0)$ $\forall i > 0\, \forall \theta$. Since Θ_0 is a singleton and $v_0 \equiv 0$ the condition simplifies to the following: A profitable mechanism exists if and only there do not exist functions $u_1(\theta_1), \dots, u_n(\theta_n)$ and a constant c such that $v_i(\theta) = u_i(\theta_i) - c$ $\forall i > 0\, \forall \theta$. But the latter simply means that there does not exist an agent $i > 0$ such that $v_i(\theta_i, \theta_{-i})$ is not constant in θ_{-i}. □

References

Aumann, Robert J. 1974. "Subjectivity and correlation in randomized strategies." *Journal of Mathematical Economics* 1 (1): 67–96. [66]

Aumann, Robert J. 1987. "Correlated equilibrium as an expression of Bayesian rationality." *Econometrica,* 1–18. [66]

Becker, Gary S. 1973. "A theory of marriage: Part I." *Journal of Political Economy* 81 (4): 813–46. [64]

Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman. 2014. "Optimal allocation with costly verification." *American Economic Review* 104 (12): 3779–813. [66]

Bergemann, Dirk, and Stephen Morris. 2016. "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics* 11 (2): 487–522. [66]

Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability,* 93–102. University of California Press. [66]

Blackwell, David. 1953. "Equivalent comparisons of experiments." *Annals of Mathematical Statistics,* 265–72. [66]

Börgers, Tilman, Angel Hernando-Veciana, and Daniel Krähmer. 2013. "When are signals complements or substitutes?" *Journal of Economic Theory* 148 (1): 165–95. [66]

Börgers, Tilman, and Peter Postl. 2009. "Efficient compromising." *Journal of Economic Theory* 144 (5): 2057–76. [66]

Brualdi, Richard A. 2006. *Combinatorial matrix classes.* Vol. 13. Cambridge University Press. [63]

Crémer, Jacques, and Richard P. McLean. 1985. "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent." *Econometrica* 53: 345–61. [53, 58, 59, 66]

Crémer, Jacques, and Richard P. McLean. 1988. "Full extraction of the surplus in Bayesian and dominant strategy auctions." *Econometrica,* 1247–57. [53, 58, 66]

Feng, Tangren, and Qinggong Wu. 2019. "Getting Information from Your Enemies." [66]

Gershkov, Alex, Jacob K. Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi. 2013. "On the equivalence of Bayesian and dominant strategy implementation." *Econometrica* 81 (1): 197–220. [67]

Goldlücke, Susanne, and Thomas Tröger. 2020. "The multiple-volunteers principle." *Available at SSRN 3753985.* [66]

Hardy, Godfrey Harold, John Edensor Littlewood, and George Pólya. 1952. *Inequalities.* Cambridge university press. [64]

Hoffman, Alan J. 1963. "On simple linear programming problems." In *Proceedings of Symposia in Pure Mathematics,* 7: 317–27. [64]

Kattwinkel, Deniz. 2020. "Allocation with Correlated Information: Too good to be true." In *Proceedings of the 21st ACM Conference on Economics and Computation,* 109–10. [66]

Kattwinkel, Deniz, and Jan Knoepfle. 2019. "Costless Information and Costly Verification: A Case for Transparency." *Available at SSRN 3426817.* [66]

Kim, Semin. 2017. "Ordinal versus cardinal voting rules: A mechanism design approach." *Games and Economic Behavior* 104: 350–71. [66]

Manelli, Alejandro M, and Daniel R Vincent. 2010. "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model." *Econometrica* 78 (6): 1905–38. [67]

Myerson, Roger B. 1981. "Optimal auction design." *Mathematics of operations research* 6 (1): 58–73. [66]

Mylovanov, Tymofiy, and Andriy Zapechelnyuk. 2017. “Optimal allocation with ex post verification and limited penalties.” *American Economic Review* 107 (9): 2666–94. [66]

Niemeyer, Axel, and Justus Preusser. 2022. “Simple Allocation with Correlated Types.” https://jpreusser.github.io/. [66]

Rosenthal, Robert W. 1974. “Correlated equilibria in some classes of two-person games.” *International Journal of Game Theory* 3 (3): 119–28. [66]

Villani, Cédric. 2009. *Optimal transport: old and new.* Vol. 338. Springer. [62]

Vince, A. 1990. “A rearrangement inequality and the permutahedron.” *American Mathematical Monthly* 97 (4): 319–23. [64]

Chapter 3

Transparency in Sequential Common-Value Trade

3.1 Introduction

In many markets, participants learn over time. A buyer who arrives to the market at a late date may learn payoff-relevant information from the fact that earlier buyers bargained with the seller but chose not to buy. The precise inference depends on the information that is made public about past interactions: how many earlier buyers bargained with the seller, and what terms of trade were they offered? We analyze the welfare implications of this kind of transparency in a dynamic market for a common-value good.

Existing work focuses on situations where the seller and no one else is initially informed about the good's value, and where the short-lived buyers propose the terms of trade (for example, Hörner and Vieille, 2009; Fuchs, Öry, and Skrzypacz, 2016; Kim, 2017). Yet, it is also plausible that that information is revealed only gradually to all market participants, and that the seller proposes the terms of trade. Why might we expect the effect of transparency to depend on these details of the market? When the seller proposes, they can affect the flow of information to the market, which is an idea that goes back at least to Taylor (1999). Indeed, a buyer's rejecting a low price is more informative than rejecting a high price. The market's transparency affects what later buyers learn from these rejections, and, therefore, the seller's incentives for charging high prices in early periods. We take a step towards understanding this interaction between learning and transparency.

In our model, there is a long-lived seller and a sequence of short-lived buyers. The seller has a single indivisible good for which the buyers have a common value. The value takes one of two values. The seller solicits buyers one-by-one. While the seller is uninformed about the value, the buyers observe informative signals from a discrete signal structure. Conditional on the value, these signals are independent and identically distributed.

In each period, the seller makes a recommendation that identifies for which signal realizations a buyer should buy the object. An intermediary then picks a price that implements the seller's recommendation. We mainly think of this formulation of the trading procedure as a convenient modelling choice that permits us to focus on the flow of information to the market; we give an interpretation once we have presented the model.

We compare three transparency regimes:

(1) The seller's past recommendations and time-on-the-market are both observable.
(2) Past recommendations are unobservable, but time-on-the-market is observable.
(3) Past recommendations and time-on-the-market are both unobservable.

For each regime, we ask whether the seller extracts the full surplus from trade. There are commonly known gains from trade, meaning that the object is traded with certainty in all equilibria. We can thus focus on the way the surplus is divided between the players.

Our main insight is that, in a sense to be made precise momentarily, the seller's ability to extract the full surplus is smallest in the intermediate regime (2). The seller would benefit from passing to either of the two more extreme regimes (1) and (3).

In all regimes, the seller extracts the full surplus if and only if buyers accrue zero information rents. This, in turn, happens if and only if trade is certain to take place with a buyer observing the most optimistic signal. When recommendations and time-on-the-market are both observable, we establish as a benchmark that the seller indeed extracts the full surplus in the unique perfect Bayesian equilibrium. It is as if the seller had commitment power.

Our first main result concerns the game with unobservable past recommendations but observable time-on-the-market. We show that, if the private signals of buyers are sufficiently rich, then the seller's utility is bounded away from the full surplus across all perfect Bayesian equilibria. By sufficiently rich we mean that for each value realization the conditional signal distribution approximates a continuous strictly positive density.

The argument is as follows: If the seller deviates (from a candidate equilibrium strategy) when meeting the first buyer, all later buyers fail to account for this deviation in their beliefs. For certain deviations, later buyers' beliefs will be too optimistic about the value relative to the correct Bayesian posterior—they are fooled into overpaying. The downside for the seller from the deviation is that if the first buyer trades, then this buyer accrues information rents. We show that this downside is dominated in rich signal structures. Using such a deviation, we argue that the unique strategy profile that would let the seller appropriate the full surplus cannot be sustained in equilibrium. (However, an equilibrium exists.)

For our second main result, we turn to the third regime where neither past recommendations nor time-on-the-market are observable. A deviation (from a candidate equilibrium strategy) can now signal that the seller has failed to trade with many buyers. Failing to trade is an indicator of poor value. Thus a deviation entails trading at terms that are quite unfavorable for the seller. Building on this intuition, we show that the seller extracts the full surplus in a sequential equilibrium. The same idea can be used construct sequential equilibria where buyers are left some surplus. In these equilibria (that may or may not leave surplus to buyers), the seller makes a constant recommendation for many early periods. The price is constant across these periods, and the good will trade with overwhelming probability at this constant price. We derive these results in a modified game where the number of buyers is finite but large, and where the seller incurs small costs for soliciting new buyers. Therefore, these sequential equilibria are sustained even in the presence of (small) incentives for trading quickly.

Since our result on the failure of full surplus extraction in the second regime assumes that signals are rich, a natural follow-up is to ask whether the seller would benefit from coarser signal structures. This is indeed the case, in the following sense: if buyers' signals about the value are binary, then buyer surplus is zero in all sequential equilibria of all three regimes.

The paper is organized as follows. In Section 3.2 we study the model with unobservable recommendations and observable time-on-the-market. The regime where everything is observable is presented as a benchmark in this section. In Section 3.3, we consider the game where neither the seller's recommendations nor time-on-the-market are observable. Section 3.4 discusses the literature, and Section 3.5 concludes. All proofs are in the appendices.

3.2 Observable time-on-the-market

3.2.1 Model

We consider a game between a seller, and countably infinitely-many buyers and intermediaries. The seller is long-lived. All other players are short-lived and arrive to the market in a pre-determined order.

3.2.1.1 Environment

The seller (she) owns a single indivisible good which she values at 0. Buyers have a common value for the good that depends on an unobservable state. The state has two possible realizations, ℓ and h, with associated values v_ℓ and v_h. We assume $0 < v_\ell < v_h$, and so it is common knowledge that there are gains from trade. Let $\alpha_{\omega,0} \in (0,1)$ be the common prior that the state is $\omega \in \{\ell, h\}$. It will frequently be more convenient to represent beliefs via the likelihood ratio of state h vs. ℓ. We denote the prior likelihood ratio by $\pi_0 = \alpha_{h,0}/\alpha_{\ell,0}$.[1]

At the start of the game, the seller is uninformed about the state. Each buyer (he) is endowed with a private signal from a finite set S. Conditional on state ω, the signals of (each finite subset of) the buyers are independent draws from a distribution f_ω that has support S.

Since the state is binary, it is without loss to order signals according to their likelihood ratios; that is, we assume

$$\forall_{s,s' \in S} \quad s < s' \Rightarrow \frac{f_h(s)}{f_\ell(s)} < \frac{f_h(s')}{f_\ell(s')}. \tag{MLRP}$$

Given $s \in S$ and $\pi \in (0, \infty)$, let

$$\hat{v}(s, \pi) = v_\ell + (v_h - v_\ell)\frac{\pi \frac{f_h(s)}{f_\ell(s)}}{\pi \frac{f_h(s)}{f_\ell(s)} + 1}. \tag{3.1}$$

The value is $\hat{v}(s, \pi)$ is the posterior value for a buyer who observes a signal realization s starting at a belief π. The prior value of the good, termed the *full surplus*, is denoted $\hat{v}_0$ and given by

$$\hat{v}_0 = v_\ell + (v_h - v_\ell)\frac{\pi_0}{\pi_0 + 1}. \tag{3.2}$$

3.2.1.2 Trading protocol

The game unfolds in discrete time, indexed by $\mathbb{N} = \{1, 2, \ldots\}$. In period i, the seller, buyer i, and intermediary i are active. First, the seller picks an element σ_i of S. We shall think of this as the seller recommending that buyers with a signal above σ_i buy the object. Accordingly, we refer to this as the recommended cutoff. The recommendation is observed by intermediary i, who then posts a price p_i. Next, buyer i arrives to the market with probability $\lambda \in (0, 1)$. Whether buyer i arrives is unobserved by all other players. If he arrives, he learns σ_i, p_i, and the realization of his private signal. He then decides whether to buy at p_i. If he does, the object is traded and the game ends. If buyer i does not arrive to the market or does not trade, the game moves to the next period.

Note that the seller's time-on-the-market is observable to buyers in the sense that each buyer i is in the market in period i or not at all.

The seller's payoff is the price at which the object is traded, if at all. Buyer i's payoff in state ω is $v_\omega - p_i$ if he trades; else his payoff is 0. As for the intermediaries

1. So, a belief of 0 means that the state is sure to be ℓ. A belief of ∞ means the state is sure to be h. All relevant Bayesian posteriors in our model will lead to beliefs in $(0, \infty)$.

we assume the following: If the seller recommends σ_i and buyer i arrives with a signal s_i such that $\sigma_i \le s_i$ but buyer i ends up not buying the good, then the payoff of intermediary i is $-\infty$. In all other cases, the intermediary's payoff is p_i. (We interpret the intermediaries further below.) The solution concept is perfect Bayesian equilibrium.

3.2.1.3 Equilibrium prices

Buyer i's beliefs about the state depend on his private signal s_i, the seller's recommendation σ_i, and the fact that he finds the good to not have been sold in previous periods.[2] Let $\pi_i(\sigma_i)$ denote i's belief (expressed as the likelihood ratio of h vs. ℓ) after learning that the game has reached round i and learning the seller's action σ_i, but before learning his private signal s_i. We will refer to π simply as buyers' beliefs. This is a slight abuse of language as π does not include a buyer's inference from his private signal, but no confusion should arise.

Once buyer i learns s_i, his valuation for the good updates to $\hat{v}(s_i, \pi_i(\sigma_i))$. Thus he is willing to accept a price p_i if and only if

$$\hat{v}(s_i, \pi_i(\sigma_i)) \ge p_i.$$

The MLRP implies that $\hat{v}(s_i, \pi_i(\sigma_i))$ is increasing in s_i. In equilibrium, the intermediary, acting according to his preferences, sets the price as large as possible subject to the constraint that buyer i accepts if s_i is weakly above σ_i, Hence the intermediary sets $p_i = \hat{v}(\sigma_i, \pi_i(\sigma_i))$ whenever the principal recommends σ_i. Buyer i will accept after a recommendation of σ_i if and only if i's private signal is weakly above σ_i.

These observations let us simplify the description of equilibrium: It suffices to specify the recommendations of the seller, and buyers' beliefs π.

The seller's conditions her recommendation only on calendar time and her past recommendations. Her set of pure strategies is thus the set S^∞ of sequences in S. Her set of mixed strategies is the set $\Delta(S^\infty)$ of distributions over S^∞.[3] Generic elements of S^∞ and $\Delta(S^\infty)$, respectively, are denoted σ and μ, respectively. Buyers' beliefs are given by a function $\pi\colon \mathbb{N} \times S \to [0, \infty]$. Given (μ, π), we denote the seller's utility by $V(\mu, \pi)$. (See the appendix for general formulas for the seller's expected utility and buyers' posteriors.)

We consider the following notion of equilibrium.

Definition 3.1. A pair (μ, π) is an *equilibrium* if μ maximizes $V(\cdot, \pi)$ across $\Delta(S^\infty)$, and π satisfies all of the following:

2. The belief could in principle also depend on the intermediary's price. Note however, that the buyer knows as much as the intermediary about the history, and hence we can safely omit this dependence.

3. The finite set S has the discrete metric, and S^∞ has the product metric. A distribution over S^∞ means a Borel probability-measure on S^∞.

- For all $s \in S$, we have $\pi_1(s) = \pi_0$.
- For all $i \geq 2$ and all $s \in S$, if s is played by μ with non-zero probability in period i (meaning $\mu(\{\sigma \in S^\infty : \sigma_i = s\}) > 0$), then $\pi_i(s)$ is derived from μ via Bayes' rule.

Some clarifying remarks are in order.

(1) All periods i are reached with non-zero probability since in all earlier periods buyers may fail to arrive. In particular, the seller never finds herself in a period off the path of play. Hence equilibrium only requires her strategy to maximize her ex-ante utility $V(\cdot, \pi)$.

(2) The first condition on the buyers' beliefs requires that the seller cannot signal what she does not: buyer 1, who is the first to interact with the seller, draws no inference from the seller's recommended cutoff in period 1 as the seller is initially uninformed about the state.

(3) The second condition on the buyers' beliefs states that all other beliefs are derived from Bayes' rule where possible: all periods i are reached with non-zero probability, and hence $\pi_i(s)$ can be derived from Bayes' rule if and only if μ plays s with non-zero probability in period i.

3.2.1.4 Sequential equilibria

Some of our results concern sequential equilibria.[4]

Definition 3.2. A strategy of the seller is *fully mixed* if in all periods it recommends all cutoffs with non-zero probability.

An equilibrium (μ, π) is a *sequential equilibrium* if there is a sequence $\{\mu_k\}_{k\in\mathbb{N}}$ of fully mixed strategies and a sequence $\{\pi_k\}_{k\in\mathbb{N}}$ of beliefs satisfying both of the following:

(1) The sequence $\{\mu_k\}_{k\in\mathbb{N}}$ weak-$*$ converges to μ, and the sequence $\{\pi_k\}_{k\in\mathbb{N}}$ converges to π pointwise.

(2) For all k, the beliefs π_k are derived from μ_k via Bayes' rule.

Lemma 3.1. *There exists a sequential equilibrium.*

3.2.1.5 Interpreting the intermediaries

We think of the seller as recommending the cutoff above which a buyer should buy the good. In order to actually implement this cutoff via a price, the seller relies on intermediaries. This could be because intermediaries have greater expertise than

4. The definition is at a slight abuse of language as we do not consider perturbations of the buyers' or intermediaries' strategies.

the seller for interacting with buyers. For concreteness, suppose the seller writes a contract that rewards the intermediary with a share ρ of the price, provided the price is low enough to implement the seller's recommendation. If there are multiple intermediaries in each period, they compete the share ρ down to 0.

That said, we mostly view the intermediaries as a convenient modelling tool. The advantage of our formulation is that we can focus on the cutoff at which trade happens. In each period, the cutoff determines the information rents that the present buyer accrues in the event of trade, and it determines what the market learns about the value in the event of no-trade. Hence the cutoffs are key to determining the division of surplus.

3.2.2 The full surplus is an upper bound

We begin our analysis by showing that the full surplus is an upper bound on the seller's equilibrium expected utility. For expositional purposes, let us assume $\lambda = 1$, meaning that buyers are sure to arrive to the market (but the results are stated for arbitrary $\lambda \in (0, 1)$). For $\lambda = 1$, the game reaches period i if and only if all preceding buyers had signals strictly below the seller's cutoff. Let $\tilde{s}_i$ denote buyer i's random signal, and let $\tilde{\sigma}_i$ denote the (possibly random) cutoff of the seller in period i. We therefore identify the event

$$\{\tilde{s}_1 < \tilde{\sigma}_1, \ldots, \tilde{s}_{i-1} < \tilde{\sigma}_{i-1}, \tilde{s}_i \geq \tilde{\sigma}_i\}$$

with the event that buyer i ends up buying the object.

Consider an equilibrium (μ, π). The seller's strategy μ together with the distribution of states and signals induces some joint distribution of recommended cutoffs, states, and signals. We denote the probability- and expectation-operators with respect to this distribution by $\mathbb{P}$ and $\mathbb{E}$.

Recall that if buyer i buys at a cutoff σ_i, he will pay $\hat{v}_i(\sigma_i, \pi_i(\sigma_i))$. This is his posterior valuation conditional on the game reaching period i, the seller recommending σ_i, and his signal being *equal* to σ_i; let us denote this valuation by

$$\mathbb{E}[v|\tilde{s}_1 < \tilde{\sigma}_1, \ldots, \tilde{s}_{i-1} < \tilde{\sigma}_{i-1}, \tilde{s}_i = \sigma_i]$$

The seller's equilibrium expected utility is therefore

$$\sum_{i=1}^{\infty} \Bigg(\mathbb{P}\left(\tilde{s}_1 < \tilde{\sigma}_1, \ldots, \tilde{s}_{i-1} < \tilde{\sigma}_{i-1}, \tilde{s}_i \geq \tilde{\sigma}_i\right) \\ \times \mathbb{E}[v|\tilde{s}_1 < \tilde{\sigma}_1, \ldots, \tilde{s}_{i-1} < \tilde{\sigma}_{i-1}, \tilde{s}_i = \tilde{\sigma}_i] \Bigg). \tag{3.3}$$

The MLRP implies that this is no greater than

$$\begin{aligned}\sum_{i=1}^{\infty}\Big(&\mathbb{P}\left(\tilde{s}_1 < \tilde{\sigma}_1,\ldots,\tilde{s}_{i-1} < \tilde{\sigma}_{i-1},\tilde{s}_i \geq \tilde{\sigma}_i\right)\\ &\times \mathbb{E}\left[v|\tilde{s}_1 < \tilde{\sigma}_1,\ldots,\tilde{s}_{i-1} < \tilde{\sigma}_{i-1},\tilde{s}_i \geq \tilde{\sigma}_i\right]\Big).\end{aligned} \tag{3.4}$$

By iterated expectations, the sum in (3.4) is nothing but the prior value of the good, namely the full surplus $\hat{v}_0$. In fact, the MLRP implies that (3.3) is strictly less than (3.4) if, with non-zero μ-probability, a period is reached where the cutoff is strictly below the largest signal in S. That is, the seller leaves information rents unless she is certain to trade with the highest possible signal.

To state this formally, let us denote by $\bar{s}$ the largest signal in S. Let $\bar{\sigma}$ be the sequence of cutoffs that is constantly equal to $\bar{s}$.

Lemma 3.2. *In all equilibria, the seller's utility is at most $\hat{v}_0$. If in an equilibirum the seller's utility is $\hat{v}_0$, then in this equilibrium the seller's strategy is $\bar{\sigma}$.*

Note that even if the seller's recommendation is constantly equal to $\bar{s}$, the price is strictly decreasing over time. Indeed, in this case the price in period i is given by

$$\mathbb{E}\left[v|\tilde{s}_1 < \bar{s},\ldots,\tilde{s}_{i-1} < \bar{s},\tilde{s}_i = \bar{s}\right].$$

As i increases, this expectation conditions on a larger number of signals being below $\bar{s}$, which depresses beliefs.[5] Put differently, as the good is not being sold, the intermediaries are decreasing prices at just the right rate to keep buyer with signals equal to $\bar{s}$ indifferent between buying.

Lemma 3.2 does not say that the strategy $\bar{\sigma}$ is actually sustained in equilibrium. Before investigating whether this can happen, let us make good on discussing the promised benchmark where past recommendation are observable.

3.2.3 Full surplus extraction with observable recommendations

Suppose for a moment that the seller's recommended cutoffs were observable. We claim that in this case she can extract the full surplus in equilibrium by playing the pure strategy $\bar{\sigma}$. In fact, this is the only equilibrium. To see this, suppose the seller uses some pure strategy σ.[6] Since the seller's actions are observable, buyer i makes

5. Precisely, the conditional expectation reads

$$v_\ell + (v_h - v_\ell)\frac{\pi_0\frac{f_h(\bar{s})}{f_\ell(\bar{s})}\left(\frac{1-\lambda f_h(\bar{s})}{1-\lambda f_\ell(\bar{s})}\right)^{i-1}}{\pi_0\frac{f_h(\bar{s})}{f_\ell(\bar{s})}\left(\frac{1-\lambda f_h(\bar{s})}{1-\lambda f_\ell(\bar{s})}\right)^{i-1}+1}.$$

which, by the MLRP, is strictly decreasing in i.

6. Since her actions are observable, pure strategies are without loss.

the correct inference from play; that is, his belief agrees with the Bayesian posterior induced by σ. The price which buyer i is offered must therefore correctly account for the Bayesian inference from reaching period i. Since this is true for all periods i, the seller's utility from σ is (assuming $\lambda = 1$)

$$\sum_{i=1}^{\infty} \Big(\mathbb{P}\left(\tilde{s}_1 < \sigma_1, \ldots, \tilde{s}_{i-1} < \sigma_{i-1}, \tilde{s}_i \geq \sigma_i\right) \\ \times \mathbb{E}\left[v | \tilde{s}_1 < \sigma_1, \ldots, \tilde{s}_{i-1} < \sigma_{i-1}, \tilde{s}_i = \sigma_i\right] \Big)$$

The arguments from the previous section imply that, if $\sigma \neq \bar{\sigma}$, then this utility is strictly less than $\hat{v}_0$, and hence strictly less than the utility from $\bar{\sigma}$.

The argument from the previous paragraph does not apply in the game with unobservable recommendations since, following a deviation in some period, later buyers do not revise their beliefs. We next show the seller may profitably deviate from $\bar{\sigma}$ by exploiting these incorrect beliefs and obtain a utility strictly above the prior value $\hat{v}_0$.

3.2.4 No full surplus with rich signals

In this section, we show that if signals are sufficiently rich, then the seller cannot extract the full surplus in equilibrium. We first make precise what we mean by a rich signal structure. Fixing a pair of cdfs (G_h, G_ℓ) on $[0, 1]$ and an integer k, consider S_k, $f_{\ell,k}$ and $f_{h,k}$ defined as follows:

$$\begin{aligned} S_k &= \left\{0, \frac{1}{k}, \ldots, 1 - \frac{1}{k}\right\} \\ \forall_{s \in S_k}, \quad f_{\omega,k}(s) &= G_\omega\left(s + \frac{1}{k}\right) - G_\omega(s). \end{aligned}$$

We say the sequence $\{(S_k, f_{h,k}, f_{\ell,k})\}_{k \in \mathbb{N}}$ *converges* to (G_h, G_ℓ).

Our result asserts that if (G_h, G_ℓ) admit well-behaved densities, then, fixing a signal structure far enough along the sequence $\{(S_k, f_{h,k}, f_{\ell,k})\}_{k \in \mathbb{N}}$, the seller cannot extract the full surplus. Note that the full surplus $\hat{v}_0 = v_\ell + (v_h - v_\ell)\frac{\pi_0}{\pi_0 + 1}$ does not depend on the signal structure.

Proposition 3.3. *Let (G_h, G_ℓ) be a pair of cdfs on $[0, 1]$. Let $\{(S_k, f_{h,k}, f_{\ell,k})\}_{k \in \mathbb{N}}$ be a sequence of finite signal structures converging to (G_h, G_ℓ).*

If G_h and G_ℓ admit continuous and strictly positive densities g_h and g_ℓ on $[0, 1]$ such that $\frac{g_h}{g_\ell}$ is strictly increasing, then the following holds for all except finitely many k: If the signal structure is given by $(S_k, f_{h,k}, f_{\ell,k})$, then the seller's utility is bounded away from $\hat{v}_0$ across all equilibria.

For the proof, it suffices to show that the pure strategy that always plays the largest signal fails to be an equilibrium (Lemma 3.2). For expositional purposes, let $\lambda = 1$. Suppressing the dependence on k, let $\bar{s} = \bar{s}_k$ denote the largest signal in the k'th signal structure.

Suppose towards a contradiction that there is an equilibrium where the seller's utility is $\hat{v}_0$ and she plays $\bar{s}$ in all periods. We consider a one-time deviation in period 1 to a cutoff strictly below $\bar{s}$. Let $s^\circ = s_k^\circ$ denote this cutoff (where we again suppress the dependence on k). The deviation will have two effects, with opposing implications for the seller's utility. The upside from the deviation is that if buyer 1 does not end up trading, then all later buyers will hold incorrect beliefs. In particular, since $s^\circ < \bar{s}$, rejecting a cutoff of s° is a *stronger* signal in favor of the *bad* state ℓ than rejecting $\bar{s}$. Therefore, all later buyers will hold a belief that is too optimistic; their willingness to pay will be too high relative to the true Bayesian posterior. The downside from the deviation is that if buyer 1 has a private signal strictly above s°, he will trade and accrue information rents.

Let us spell this out in more detail. Since buyer 1's belief does not react to the seller's action in round 1, the contribution from buyer 1 to the utility from the deviation is

$$\mathbb{P}(\tilde{s}_1 \geq s^\circ)\mathbb{E}[v|\tilde{s}_1 = s^\circ].$$

Now consider buyer $i > 1$. The probability that he will trade under the deviation is

$$\mathbb{P}\left(\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \ldots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right)$$

Buyer i receives his on-path recommendation, and hence his belief equals his on-path belief. Since the candidate equilibrium has the seller recommend $\bar{s}$ in all periods, the price that buyer i pays, if he trades, is

$$\mathbb{E}\left[v|\tilde{s}_1 < \bar{s}, \tilde{s}_2 < \bar{s}, \ldots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right].$$

The seller's utility from the deviation is therefore

$$\begin{aligned}
&\mathbb{P}(\tilde{s}_1 \geq s^\circ)\mathbb{E}[v|\tilde{s}_1 = s^\circ] \\
&+\sum_{i=2}^{\infty}\Bigg(\mathbb{P}\left(\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \ldots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right) \\
&\qquad\qquad \times \mathbb{E}\left[v|\tilde{s}_1 < \bar{s}, \tilde{s}_2 < \bar{s}, \ldots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]\Bigg).
\end{aligned}$$

Let us compare this to the full surplus $\hat{v}_0$ (which is the seller's utility from constantly recommending $\bar{s}$). By iterated expectations, we may write $\hat{v}_0$ as

$$\begin{aligned}&\mathbb{P}(\tilde{s}_1 \geq s^\circ)\mathbb{E}[v|\tilde{s}_1 \geq s^\circ]\\&+\sum_{i=2}^{\infty}\Big(\mathbb{P}\big(\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\big)\\&\qquad\times \mathbb{E}\left[v|\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]\Big).\end{aligned}$$

Thus the deviation is profitable if and only if

$$\begin{aligned}&\mathbb{P}(\tilde{s}_1 \geq s^\circ)\big(\mathbb{E}[v|\tilde{s}_1 = s^\circ] - \mathbb{E}[v|\tilde{s}_1 \geq s^\circ]\big)\\&+\sum_{i=2}^{\infty}\Big(\mathbb{P}\big(\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\big)\\&\qquad\times\Big(\mathbb{E}\left[v|\tilde{s}_1 < \bar{s}, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]\\&\qquad\quad - \mathbb{E}\left[v|\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]\Big)\Big)\end{aligned}$$

is strictly positive. The difference

$$\mathbb{E}[v|\tilde{s}_1 = s^\circ] - \mathbb{E}[v|\tilde{s}_1 \geq s^\circ]$$

is strictly negative, as we infer from the MLRP; this is the information rent left to buyer 1. Each term inside the infinite sum, however, is strictly positive. To see this, note that

$$\mathbb{E}\left[v|\tilde{s}_1 < \bar{s}, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]$$

conditions on $(i-1)$-many buyers failing to trade at $\bar{s}$. However,

$$\mathbb{E}\left[v|\tilde{s}_1 < s^\circ, \tilde{s}_2 < \bar{s}, \dots, \tilde{s}_{i-1} < \bar{s}, \tilde{s}_i = \bar{s}\right]$$

conditions on one buyer failing to trade at s° and $(i-2)$-many buyers failing to trade at $\bar{s}$. Since not trading at s° is a stronger signal for the bad state than not trading at $\bar{s}$, each term inside the infinite sum is strictly positive.

The seller thus benefits from the deviation if late buyers' wrong beliefs outweigh buyer 1's information rents. The proof of Proposition 3.3 shows that this happens whenever the signal structure is sufficiently rich. If we choose $s^\circ = s_k^\circ$ ever closer to 1 as $k \to \infty$, the fact that the likelihood ratio $\frac{g_h}{g_\ell}$ is continuous implies that both the loss due information rents as well as the gain due to incorrect beliefs vanish. By using that the likelihood ratio is bounded at the top (which is implied by the fact that the densities are continuous and strictly positive), we show that the loss vanishes more rapidly than the gain for a suitable choice of s_k°. In particular, this is the case when s_k° converges an order of magnitude more slowly to 1 than $\bar{s} = \bar{s}_k$.

3.2.5 Full surplus with binary signals

Since Proposition 3.3 concerns rich signals, a natural follow-up question asks how the surplus is divided when signals are coarse. The next result shows that when signals are binary we reach a conclusion starkly different from Proposition 3.3.

Proposition 3.4. *Let signals be binary, meaning $|S| = 2$. If (μ, π) is a sequential equilibrium, then the seller's strategy is the pure strategy $\bar{\sigma}$, and her expected utility is the full surplus $\hat{v}_0$.*

The key observation is that, for binary signals, no strategy induces more pessimistic beliefs than constantly playing the highest cutoff $\bar{s}$. Let $\underline{s}$ denote the smallest signal, which here simply means the only signal different from $\bar{s}$. A recommendation of $\underline{s}$ is accepted whenever a buyer arrives to the market. Failing to trade at $\underline{s}$ therefore reveals that no buyer arrived to the market. Since this event contains no information about the value, the belief remains unchanged. Conversely, whenever $\bar{s}$ does not lead to a trade, the posterior that the state is h decreases.

For a sequential equilibrium, the observation from the previous paragraph implies that the posterior induced by $\bar{\sigma}$ is also a lower bound on buyers' off-path beliefs. Hence a lower bound on equilibrium utility is given by the utility from deviating to $\bar{\sigma}$ and forming the induced prices using the beliefs induced by $\bar{\sigma}$. But this lower utility is nothing but $\hat{v}_0$, as we infer from the discussion in Section 3.2.2.[7]

3.3 Unobservable time-on-the-market

3.3.1 Model

In this section, we consider another game. Its defining property is that buyers observe neither the seller's past recommendations nor the seller's time-on-the-market. That is, relative to the game of the previous section, a buyer is now also unaware of the label of the period in which he is asked to make a move.

The number n of buyers and intermediaries is now finite. At the beginning of the game, Nature picks a permutation of $\{1, \ldots, n\}$ according to the uniform distribution. The realized permutation is not observed by any player, and it determines the order in which buyers and intermediaries arrive to the market. When asked to make a move, each buyer observes his private signal, the intermediary's price, and the seller's recommendation. Each intermediary observes the seller's current recommendation. When in period i the seller recommends a cutoff σ_i, the buyer's posterior

7. The argument sketched here uses the assumption that arrivals to the market are probabilistic, meaning $\lambda \in (0, 1)$. Suppose that arrivals are certain, $\lambda = 1$. Playing the lowest signal now means trading with probability one. In a sequential equilibrium, the beliefs of buyers who are reached with probability zero along the path of play must therefore equal the beliefs induced by $\bar{\sigma}$. A similar argument thus shows that the deviation to $\bar{\sigma}$ must still yield $\hat{v}_0$.

belief (expressed as the likelihood ratio of h vs. ℓ) is denoted $\pi^{\emptyset}(\sigma_i)$.[8] In the same situation, the intermediary will find it optimal to choose a price of $\hat{v}(\sigma_i, \pi^{\emptyset}(\sigma_i))$. As before, a buyer finds it optimal to accept σ_i at a private signal s if and only if s is weakly greater than σ_i.

On the seller's side, we now assume that she incurs a cost c whenever the game moves to the next period. This cost can be interpreted as costs for soliciting new buyers, and we assume $c \in [0, \lambda v_\ell]$.[9]

A mixed strategy of the seller is now a distribution μ over the set S^n of finite cutoff sequences. Buyers' beliefs are represented by $\pi^{\emptyset} : S \to [0, \infty]$. The seller's profit from this pair is denoted $V^{\emptyset}(\mu, \pi^{\emptyset}, n, c)$. (The appendix presents formulas for the seller's utility and buyers' Bayesian posteriors.) Let $\Gamma^{\emptyset}(n, c)$ denote the game described here.

Definition 3.3. A pair $(\mu, \pi^{\emptyset})$ is an *equilibrium* of $\Gamma^{\emptyset}(n, c)$ if μ is a maximizer of $V^{\emptyset}(\cdot, \pi^{\emptyset}, n, c)$, and $\pi^{\emptyset}$ satisfies the following: For all $s \in S$, if s is played by μ with non-zero probability in some period (meaning $\sum_{i=1}^{n} \mu(\{\sigma \in S^n : \sigma_i = s\}) > 0$), then $\pi^{\emptyset}(s)$ is derived from μ via Bayes' rule.

A strategy of the seller is *fully mixed* if for all cutoffs there is at least one period in which the cutoff is played with non-zero probability; that is, all $s \in S$ satisfy $\sum_{i=1}^{n} \mu(\{\sigma \in S^n : \sigma_i = s\}) > 0$.[10]

An equilibrium $(\mu, \pi^{\emptyset})$ is a *sequential equilibrium* if there is a sequence $\{\mu_k\}_{k \in \mathbb{N}}$ of fully mixed strategies and a sequence $\{\pi_k^{\emptyset}\}_{k \in \mathbb{N}}$ of beliefs satisfying both of the following.

(1) The sequence $\{\mu_k\}_{k \in \mathbb{N}}$ converges to μ, and the sequence $\{\pi_k^{\emptyset}\}_{k \in \mathbb{N}}$ converges to $\pi^{\emptyset}$.[11]

(2) For all k, the beliefs $\pi_k^{\emptyset}$ are derived from μ_k via Bayes' rule.

Lemma 3.5. *For all $n \in \mathbb{N}$ and $c \in [0, \lambda v_\ell]$ there exists a sequential equilibrium of $\Gamma^{\emptyset}(n, c)$.*

3.3.2 Signaling calendar time

Our aim in this section is to show that, along a certain sequence of sequential equilibria, the seller can extract the full surplus as the number of buyers grows large and solicitation costs vanish. However, not all sequences of sequential equilibria have this

8. Mnemonically, the superscript $\emptyset$ indicates that buyers know neither the seller's past actions nor her time-on-the-market.

9. The assumption that c is in $[0, \lambda v_\ell]$ implies that the seller will always find it optimal to keep searching until the pool of buyers is exhausted. Specifically, she can always recommend the lowest signal as a cutoff, leading to trade at a price of at least v_ℓ when a buyer arrives with probability λ.

10. Notice that this notion of a fully mixed strategy differs from the regime with observable time-on-the-market.

11. All strategies and beliefs are viewed as elements of Euclidean space.

property. To state the result formally, given $s \in S$, let $\bar{F}_\omega(s)$ denote the probability of observing a signal weakly above s. Recall also that $\underline{s}$ denotes the smallest signal.

Proposition 3.6. *Let $s^* \in S \setminus \{\underline{s}\}$. Let $\{c_n\}_{n\in\mathbb{N}}$ be a sequence in $[0, \lambda v_\ell]$ converging to* 0. *For all $n \in \mathbb{N}$, there exists μ_n, $\pi_n^\emptyset$, and an integer j_n such that the sequence $\{(\mu_n, \pi_n^\emptyset, j_n)\}_{n\in\mathbb{N}}$ satisfies all of the following:*

(1) For all but finitely many $n \in \mathbb{N}$, the pair $(\mu_n, \pi_n^\emptyset)$ is a sequential equilibrium of $\Gamma^\emptyset(n, c_n)$.

(2) For all $n \in \mathbb{N}$, the seller using μ_n plays s^* in the first j_n rounds with probability one; that is, we have $\mu_n(\{\sigma \in S^n : (\sigma_1, \dots, \sigma_{j_n}) = (s^*, \dots, s^*)\}) = 1$.

(3) The sequence $\{j_n\}_{n\in\mathbb{N}}$ diverges to ∞.

(4) Along the sequence, the good is traded with probability converging to 1. *The seller's expected utility and the price at which the good is traded converge almost surely to*

$$v_\ell + (v_h - v_\ell)\frac{\pi\frac{f_h(s^*)}{f_\ell(s^*)}\frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)}}{\pi\frac{f_h(s^*)}{f_\ell(s^*)}\frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} + 1}. \tag{3.5}$$

In the special case $s^* = \bar{s}$ we have $\bar{F}_\omega(\bar{s}) = f_\omega(\bar{s})$ for all ω, and hence the price in (3.5) equals $\hat{v}_0$. That is, the seller gets the full surplus along this sequence of equilibria. Whenever s^* is different from $\bar{s}$, however, the (MLRP) implies

$$\frac{f_h(s^*)}{f_\ell(s^*)}\frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} < 1.$$

Thus, for s^* different from $\bar{s}$, the seller's equilibrium utility converges to a value strictly below $\hat{v}_0$.

The basic observation that we use for the proof of Proposition 3.6 is that the seller's recommendation contains information about her time-on-the-market. Let us sketch the proof idea for the case $s^* = \bar{s}$. Suppose for a moment that for some integer j the seller uses a pure strategy σ that recommends $\bar{s}$ in all of the first j periods, and never thereafter. When the seller recommends to a buyer an on-path cutoff different from $\bar{s}$, this reveals that the seller has unsuccessfully tried to sell the object for at least j rounds. Since failing to trade the object depresses beliefs, picking a cutoff different from $\bar{s}$ thus leads to a price approximately equal to v_ℓ, provided j is sufficiently large. Let us compare to this to the price from recommending $\bar{s}$. Since $\bar{s}$ is on-path under σ, a buyer's belief $\pi^\emptyset(\bar{s}, \sigma)$ after arriving to the market and being recommended $\bar{s}$ can be computed via Bayes' rule. This belief is given by

$$\pi^\emptyset(\bar{s}, \sigma) = \pi_0 \frac{\sum_{i=1}^{n} \frac{1}{n}\mathbf{1}_{(\sigma_i = \bar{s})}(1 - \lambda f_h(\bar{s}))^{i-1}}{\sum_{i=1}^{n} \frac{1}{n}\mathbf{1}_{(\sigma_i = \bar{s})}(1 - \lambda f_\ell(\bar{s}))^{i-1}} = \pi_0 \frac{\sum_{i=1}^{j}(1 - \lambda f_h(\bar{s}))^{i-1}}{\sum_{i=1}^{j}(1 - \lambda f_\ell(\bar{s}))^{i-1}}.$$

Evaluating the geometric sums shows that, for large j, this belief approximately equals $\pi_0 \frac{f_\ell(\bar{s})}{f_h(\bar{s})}$. The price $\hat{v}(\bar{s}, \pi^\emptyset(\bar{s}, \sigma))$ after $\bar{s}$ equals the posterior valuation conditional on a private signal of $\bar{s}$ and conditional on arriving to the market and being recommended $\bar{s}$. Hence this price approximately equals

$$v_\ell + (v_h - v_\ell)\frac{\pi_0 \frac{f_\ell(\bar{s})}{f_h(\bar{s})}\frac{f_h(\bar{s})}{f_\ell(\bar{s})}}{\pi_0 \frac{f_\ell(\bar{s})}{f_h(\bar{s})}\frac{f_h(\bar{s})}{f_\ell(\bar{s})} + 1} = v_\ell + (v_h - v_\ell)\frac{\pi_0}{\pi_0 + 1}.$$

This is nothing but the full surplus $\hat{v}_0$.

In summary, a cutoff of $\bar{s}$ yields a price of $\hat{v}_0$, whereas deviations from $\bar{s}$ yields a price of v_ℓ (approximately, when j is large). This suggests that the seller's offering $\bar{s}$ for a large number periods j can actually be sustained in equilibrium, leading to trade with overwhelming probability at $\hat{v}_0$, and hence to the seller's extracting the full surplus. A complication in this argument is that the seller also incurs costs for solicitng new buyers and that the pool of buyers is finite. Since $\bar{s}$ leads to the smallest per-period probability of trade, the seller has an incentive to deviate from $\bar{s}$ to save on costs or, when the pool of buyers is almost exhausted, to ensure a last minute sale of the object. Hence we have to consider the possibility that the seller plays signals other than $\bar{s}$ along the equilibrium path. This complicates the construction of equilibrium; care has to be taken to let j (the number of initial periods in which the seller constantly recommends $\bar{s}$) diverge to ∞, but not too rapidly.

The proof for general $s^* \in S \setminus \{\underline{s}\}$ is similar. To understand why, suppose the seller's strategy is to play s^* for the first j rounds. As long as s^* is not the lowest signal $\underline{s}$, failing to trade at s^* depresses beliefs.[12] Hence the earlier reasoning implies that deviating from s^* leads to a price approximately equal to v_ℓ, while s^* leads to a strictly higher price (namely the price in (3.5)).

3.3.3 Full surplus with binary signals

Proposition 3.6 implies that the players may fail to coordinate on a seller-optimal equilibrium whenever there are at least three signals. We conclude this section by addressing the case of binary signals. In parallel to Proposition 3.4, we find that the seller extracts the full surplus along all sequences of sequential equilibria.

Proposition 3.7. *Let signals be binary, meaning $|S| = 2$. Let $\{c_n\}_{n\in\mathbb{N}}$ be a sequence in $[0, \lambda v_\ell]$ converging to 0. For all $n \in \mathbb{N}$, let $(\mu_n, \pi_n^\emptyset)$ be a sequential equilibrium of $\Gamma^\emptyset(n, c_n)$.*[13] *As $n \to \infty$, the seller's utility along the sequence of equilibria converges to $\hat{v}_0$.*

12. As remarked at an earlier point, when $\underline{s}$ does not lead to a trade in some round, the Bayesian inference is that no buyer arrived to the market in that round. Non-arrivals reveal nothing about the state of the world, and hence playing $\underline{s}$ for many rounds will not depress beliefs towards zero.

13. As we recall, Lemma 3.5 implies that $\Gamma^\emptyset(n, c_n)$ admits a sequential equilibrium for all n.

Our proof uses the same ideas as our proof of Proposition 3.4. Namely, with binary signals, the belief $\pi_n^{\emptyset}(\bar{s})$ must be bounded below by the posterior induced by $\bar{\sigma}$ (verifying this in the present game is more complicated than in the game with observable time). It is then easy to verify that the utility from deviating to $\bar{\sigma}$ admits a lower bound which converges $\hat{v}_0$. Since the utility from this deviation is itself a lower bound on equilibrium utility and since equilibrium utility is bounded above by $\hat{v}_0$, the claim follows.

3.4 Related literature

This paper is related to the literature on transparency in dynamic markets with adverse selection. See Hörner and Vieille (2009), Kaya and Liu (2015), Fuchs, Öry, and Skrzypacz (2016), Kim (2017), and Kaya and Roy (2022a, 2022b, 2022c). It is by now understood that the effect of transparency on market outcomes is dependent on other the details of the market. For example, Kaya and Roy (2022c) show that the effect varies subtly with intra-period competition between buyers and the prior belief about quality. Our contribution to this literature is to analyze the effect of transparency when the long-lived seller makes all offers, buyers have private information, and all players are initially uninformed. Existing papers in this literature study markets where the seller is informed and buyers make the offers.[14]

The paper of Kim (2017) is perhaps closest. In Kim's model, uninformed buyers make private offers to the seller of a single unit. The seller's costs and the object's value are the seller's private information. Kim compares two regimes that differ in whether buyers observe the seller's time-on-the-market. (Kim also studies a version of the model with an inflow of new sellers and buyers.) When time-on-the-market is observable, buyers update on the fact that high types of the seller are more willing to wait for favorable terms, leading buyers to offer higher prices as the game progresses. Hence transparency (of time-on-the-market) affects the seller's incentives to delay trade. In contrast, in our model, the seller benefits from high prices that minimize buyers' information rents and delay trade. Transparency affects the seller's incentives to deviate to lower prices.

In several other papers on dynamic markets with learning, the seller is initially informed about the value of the good, but additional signals arrive to the market over time. In Zhu (2012), Lauermann and Wolinsky (2016), and Kaya and Kim (2018), buyers observe private signals about the value. In Daley and Green (2012), buyers' signals are public. We differ from these papers in that, in our model, there is no initial information asymmetry, but asymmetry develops endogenously through delay in trade and buyers' private signals.[15] Note that the lack of initial asymmetry

14. Some papers, e.g. Kaya and Liu (2015), reverse the roles of sellers and buyers.

15. In this regard, we are similar to Hwang (2018), in whose model there is no initial asymmetry, but asymmetry grows as the *seller* observes an exogenous private signal.

matters: in the game with observable time and unobservable actions, we use the assumption that the first buyer does not revise his beliefs after observing the seller's deviation (Proposition 3.3). One can show that the deviation used in the proof is not profitable if we let this buyer revise his beliefs arbitrarily.

In the regime with unobservable time, the seller's action can signal calendar time, and hence the good's value. The idea that the seller's actions signal information about value is not new. For example, Barsanetti and Camargo (2022) recently explore this idea in a model where the seller is informed about the value. Lauermann and Wolinsky (2016), who study a regime with unobservable time, also discuss this idea as part of a robustness check. A key distinguishing point is again that in our model the seller is *not* initially informed about the value. Hence, in a consistent assessment, buyers' interpretations of the sellers' actions are constrained by what the seller can actually learn along the path of play. These constraints are absent when the seller is initially informed, thereby expanding the set of equilibria. In particular, our results for binary signal structures do not extend to informed-seller settings.

Taylor (1999) considers a two-period model that is related to ours. In each period, buyers bid for the good and the seller sets a reserve price. Taylor discusses the effects of the reserve price on the speed of learning: as in our benchmark model with observable actions, the seller gains from setting high initial prices to keep future beliefs high. Taylor further notes that high types of the seller benefit from public records. We instead focus on the seller's ability to extract the full surplus in a different informational setup and with a large number of buyers.

Bose et al. (2006, 2008) study a model close to ours. Namely, a version of our benchmark with observable actions but where the seller has an infinite number of units (and chooses prices, rather than recommendations). The history of prices and sales is public. Bose et al. (2006) study whether the monopolist's strategy triggers herding behavior. Bose et al. (2008) characterize optimal offers when signals are binary. Their results have no immediate counterparts in our benchmark model with observable actions as we consider the sale of a single unit. Specifically, since selling a unit is good news about the value, their model admits belief dynamics that are absent in ours.

Bose, Orosel, and Vesterlund (2002) also consider unobservable price offers when signals are binary and sales are observable. Their Lemma 10 shows that the seller may be unable to commit to trading exclusively with the most optimistic signal. The result is driven by the fact that selling a unit is good news about the value, and hence accelerating trades makes later buyers more optimistic. As noted above, this effect is absent in our model with a single unit. Indeed, their Lemma 10 sharply contrasts our results for binary signals.

There are further papers investigating other notions of transparency in more distant settings. The following are some examples. In the bilateral bargaining of Hwang and Li (2017), the focus is on transparency of on one party's outside option. In the multilateral bargaining game of Krasteva and Yildirim (2012), the focus is on transparency of the negotiation sequence and prices. Chaves (2019) studies how the transparency of on-going negotiations affects the incentives of third parties to interrupt these negotiations. Dilmé (2022) studies imperfect signals about a long-lived players actions in repeated bargaining. In the reputation models of Pei (2022a, 2022b), the question is how limited observability of the long-lived player's actions affect that players ability to build a reputation.

3.5 Conclusion

In a dynamic market for a common value good, we have uncovered a sense in which hiding information about the seller's actions but disclosing information about her time-on-the-market is beneficial to buyers. For future work, it is interesting to consider what changes if the seller has multiple objects for sale (as in the work of Bose et al. (2006, 2008)) or if there are multiple sellers whose goods have correlated values. With multiple objects, one can investigate how the transparency of sales affect equilibrium outcomes. A different intriguing direction could attempt to endogenize buyers' arrival to the market. Our results use that the (random) order in which buyers arrive to the market is exogenous. What would change if buyers could strategically time when to solicit an offer from the seller? Lastly, our results suggest interesting open questions for information design. While we have shown that binary signals are optimal for the seller in the limit game, it is open what signal structures minimize the seller's revenue. Relatedly, which signal structures would maximize or minimize the overall surplus when there are frictions?

Appendix 3.A Observable time-on-the-market

3.A.1 Definitions and notation

This part of the appendix derives the expressions for buyers' posteriors belief and the seller's expected utility.

Since S is finite, the set S^∞ of sequences in S is compact and metric (in the product metric). This renders $\Delta(S^\infty)$ a compact metrizable space (Aliprantis and Border, 2006, Theorem 15.11). Let Π denote the set of functions from $\mathbb{N} \times S$ to $[0, \pi_0]$. As a countable product of compact intervals, the set Π is a compact metric space when equipped with the product metric.

In the main text, we initially introduced buyers beliefs as functions mapping to $[0, \infty]$. As we will see, on-path beliefs always lie in $[0, \pi_0]$. Restricting off-path beliefs to $[0, \pi_0]$ does not eliminate equilibria (since the prices, and hence the seller's utility, are increasing in buyers' beliefs). Hence there is no loss in viewing beliefs as element of Π.

Let F_ω denote the cdf. of the signals in state ω. For all s in S, let us define $\underline{F}_\omega = F_\omega(s) - f_\omega(s)$ as the probability of observing a signal strictly below s. Further, let $\bar{F}_\omega(s) = 1 - \underline{F}_\omega(s)$ denote the probability of observing a signal weakly above s.

3.A.1.1 The seller's expected utility

Let $\pi \in \Pi$ and $\sigma \in S^\infty$. When the seller uses the pure strategy σ, then in state ω the game reaches period i with probability $\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_j)\big)$. Conditional on reaching period i, buyer i ends up buying the object with probability $\lambda\bar{F}_\omega(\sigma_i)$; in that case, given beliefs π, he pays $\hat{v}(\sigma_i, \pi_i(\sigma_i))$. So the seller's expected utility equals

$$V(\sigma,\pi) \;=\; \sum_{i=1}^{\infty}\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0}\lambda\bar{F}_\omega(\sigma_i)\left(\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_j)\big)\right)\hat{v}(\sigma_i,\pi_i(s)). \tag{3.A.1}$$

The infinite sum is well-defined since for all $s \in S$ we have

$$1-\lambda\bar{F}_\omega(s) \;\le\; 1-\lambda f_\ell(\bar{s}) \;<\; 1, \tag{3.A.2}$$

meaning that $\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_j)\big)$ is bounded above by $(1-\lambda f_\ell(\bar{s}))^{i-1}$.

Using the bound in (3.A.2) and finiteness of S, a routine argument shows that $V(\sigma,\pi)$ is continuous in (σ,π). Hence V is bounded. Thus it makes sense to define the seller's expected utility from a mixed strategy μ as

$$V(\mu,\pi) \;=\; \int_{\sigma\in S^\infty} V(\sigma,\pi)\,d\mu(\sigma).$$

Using the bound in (3.A.2) and finiteness of S once again, we also find that V is continuous on $\Delta(S^\infty)\times\Pi$.

3.A.1.2 Buyers' beliefs

Given $\mu \in \Delta(S^\infty)$ and an integer i, let $S(i,\mu)$ denote the set of signals $s \in S$ satisfying

$$\int_{\sigma\in S^\infty} \mathbf{1}_{(\sigma_i=s)}\,d\mu(\sigma) > 0.$$

These are the signals s which μ plays with non-zero probability in round i. For all $s \in S(i,\mu)$, let

$$\hat{\pi}_i(s,\mu) \;=\; \pi_0\frac{\int_{\sigma\in S^\infty}\mathbf{1}_{(\sigma_i=s)}\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_h(\sigma_j)\big)\,d\mu(\sigma)}{\int_{\sigma\in S^\infty}\mathbf{1}_{(\sigma_i=s)}\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\ell(\sigma_j)\big)\,d\mu(\sigma)} \tag{3.A.3}$$

denote the Bayesian posterior likelihood of h versus ℓ conditional on reaching period i and the seller then offering a cutoff of s.

It is not difficult to see that if $s \in S(i, \mu)$ and $\{\mu_k\}_k$ is a sequence that weak-$*$ converges to μ, then $s \in S(i, \mu_k)$ holds for all but finitely many k. In the same situation, the posterior $\hat{\pi}_i(s, \mu_k)$ is well-defined for all but finitely many k and converges to $\hat{\pi}_i(s, \mu)$.

3.A.2 Equilibrium Existence

Proof of Lemma 3.1. Let μ_0 denote the strategy with the property that all $i \in \mathbb{N}$ and $s \in S$ satisfy $\mu_0\{\sigma \in S^\infty : \sigma_i = s\} = \frac{1}{|S|+1}$. That is, the seller randomizes uniformly over S in each period. This strategy μ_0 exists as one may verify, say, via an application of Ionescu-Tulcea's theorem (Bogachev, 2007, Theorem 10.7.3).

Given a strategy μ' and an integer k, note that $(1 - \frac{1}{k})\mu' + \frac{1}{k}\mu_0$ is fully mixed. Hence the Bayesian posterior belief as defined (3.A.3) induced by $(1 - \frac{1}{k})\mu' + \frac{1}{k}\mu_0$ is well-defined. Let us denote this belief by $\hat{\pi}(\cdot|\mu', k)$.[16] Consider the correspondence

$$\mu' \;\mapsto\; \operatorname*{arg\,max}_{\mu \in \Delta(S^\infty)} V(\mu, \hat{\pi}(\cdot|\mu', k))$$

As observed in Appendix 3.A.1.2, buyer's beliefs are continuous in the seller's strategy when the strategy is fully mixed. That is $\mu' \mapsto \hat{\pi}(\cdot|\mu', k)$ is continuous. We further noted in Appendix 3.A.1.1 that V is jointly continuous in the seller's strategy and beliefs. An application of Berge's Maximum Theorem (Aliprantis and Border, 2006, Theorem 17.31) implies that the above arg max-correspondence is non-empty and compact-valued, and upper-hemicontinuous. Since V is linear in the seller's strategy, the arg max-correpondence is convex-valued, too. We thus infer from the Kakutani-Fan-Glicksberg Theorem (see e.g. Corollary 17.55 of Aliprantis and Border (2006, p. 583)) that for all k there exists a strategy μ_k^* satisfying

$$\mu_k^* \;\in\; \operatorname*{arg\,max}_{\mu \in \Delta(S^\infty)} V(\mu, \hat{\pi}(\cdot|\mu_k^*, k)).$$

Let π_k^* denote the belief $\hat{\pi}(\cdot|\mu_k^*, k)$.

By compactness of $\Delta(\Sigma^\infty)$ and Π, the sequence $\{\mu_k^*, \pi_k^*\}_{k \in \mathbb{N}}$ admits a convergent subsequence. Let this be the sequence itself, and let (μ^*, π^*) denote the limit. We claim that (μ^*, π^*) is a sequential equilibrium. To that end, we note that μ^* is the limit of the sequence

$$\left\{\left(1 - \frac{1}{k}\right)\mu_k^* + \frac{1}{k}\mu_0\right\}_{k \in \mathbb{N}}.$$

16. That is, $\hat{\pi}(\cdot|\mu', k)$ is defined for all i and s by

$$\hat{\pi}_i(s|\mu', k) \;=\; \hat{\pi}_i\left(s, \left(1 - \frac{1}{k}\right)\mu' + \frac{1}{k}\mu_0)\right).$$

For all k, the strategy $\left(1-\frac{1}{k}\right)\mu_k^* + \frac{1}{k}\mu_0$ is fully mixed and the belief π_k^* is obtained from $\left(1-\frac{1}{k}\right)\mu_k^* + \frac{1}{k}\mu_0$ via Bayes' rule. Therefore, to show that (μ^*, π^*) is a sequential equilibrium, it suffices to show that μ^* maximizes $V(\cdot, \pi^*)$ across $\Delta(S^\infty)$. Letting μ be an arbitrary strategy, we know that for all k we have $V(\mu_k^*, \pi_k^*) \geq V(\mu, \pi_k^*)$. Taking $k \to \infty$ and using continuity of V, we infer that $V(\mu^*, \pi^*) \geq V(\mu, \pi^*)$ holds, as promised.

□

3.A.3 Failure of surplus extraction

3.A.3.1 Auxiliary results

Proof of Lemma 3.2. Let (μ, π) be an equilibrium. Let $i \in \mathbb{N}$ and $\sigma_i \in S(i, \mu)$. If trade happens at (i, σ_i), the price equals

$$\hat{v}(\sigma_i, \pi_i(\sigma_i)) \;=\; v_\ell + (v_h - v_\ell)\frac{\pi_i(\sigma_i)\frac{f_h(\sigma_i)}{f_\ell(\sigma_i)}}{\pi_i(\sigma_i)\frac{f_h(\sigma_i)}{f_\ell(\sigma_i)} + 1}.$$

Since $\sigma_i \in S(i, \mu)$, the belief $\pi_i(\sigma_i)$ is derived from Bayes' rule. Using (3.A.3), one may verify that $\pi_i(\sigma_i) \in (0, \infty)$ holds. By the MLRP, the price in the previous display is no greater than

$$v_\ell + (v_h - v_\ell)\frac{\pi_i(\sigma_i)\frac{\bar{F}_h(\sigma_i)}{\bar{F}_\ell(\sigma_i)}}{\pi_i(\sigma_i)\frac{\bar{F}_h(\sigma_i)}{\bar{F}_\ell(\sigma_i)} + 1}.$$

This is the posterior valuation conditional on the joint event that a signal above σ_i realizes, the game reaches period i, and the sellers recommends σ_i; let us denote this event by $E_i(\sigma_i)$. The posterior valuation conditional on $E_i(\sigma_i)$ is $\mathbb{E}[v|E_i(\sigma_i)]$. Note that since $\pi_i(\sigma_i) \in (0, \infty)$ we have $\hat{v}(\sigma_i, \pi_i(\sigma_i)) < \mathbb{E}[v|E_i(\sigma_i)]$ whenever $\sigma_i < \bar{s}$ holds.

Trade happens at (i, σ_i) if and only if the event $E_i(\sigma_i)$ occurs. We know from the bound in (3.A.2) that the probability of not trading within the first i rounds converges to 0 as $i \to \infty$, uniformly across all strategies of the seller. Put differently, as $i \to \infty$, the probability that the event

$$\bigcup_{(\sigma_1,\ldots,\sigma_i)\in S^i}\left(\bigcup_{j=1}^{i} E_j(\sigma_j)\right)$$

does not occur converges to 0. It follows from the Law of Iterated Expectations that the seller's profit is at most the prior valuation $\hat{v}_0$, with equality if and only if the induced cutoff in each period is $\bar{s}$ with probability one. The unique strategy for which this can hold is therefore the pure strategy $\bar{\sigma}$. □

Lemma 3.8. *Let (g_h, g_ℓ) and $\{(S_k, f_{h,k}, f_{\ell,k})\}_{k\in\mathbb{N}}$ be as in the hypothesis of Proposition 3.3. For all k, let $\bar{s}_k = 1 - 1/k$. There exists a sequence $\{s_k^\circ\}_{k\in\mathbb{N}}$ such that for all except finitely-many k we have $s_k^\circ \in S_k$, and such that (the following limits exist and satisfy)*

$$\infty > \lim_{k\to\infty} \frac{f_{h,k}(\bar{s}_k)}{f_{\ell,k}(\bar{s}_k)} > 1, \tag{3.A.4a}$$

$$\lim_{k\to\infty} \frac{f_{h,k}(\bar{s}_k)}{1 - \underline{F}_{\ell,k}(s_k^\circ)} = 0, \tag{3.A.4b}$$

$$\lim_{k\to\infty} \frac{f_{h,k}(\bar{s}_k)}{f_{\ell,k}(\bar{s}_k)} \frac{f_{\ell,k}(s_k^\circ)}{f_{h,k}(s_k^\circ)} = 1. \tag{3.A.4c}$$

Proof of Lemma 3.8. For all k, let $s_k^\circ = \max\{s \in S_k : s \le 1 - 1/\sqrt{k}\}$. Note that we have $\bar{s}_k = 1 - 1/k$.

Considering (3.A.4a), we note that $\frac{f_{h,k}(\bar{s}_k)}{f_{\ell,k}(\bar{s}_k)} = \frac{1-G_h(1-1/k)}{1-G_\ell(1-1/k)}$ converges to $g_h(1)/g_\ell(1)$, as an application of L'Hôpital's rule shows. This limit is strictly greater than one.

Next consider (3.A.4b). The ratio $\frac{f_{h,k}(\bar{s}_k)}{1-\underline{F}_{\ell,k}(s_k^\circ)}$ equals $\frac{1-G_h(1-1/k)}{1-G_\ell(1-1/\sqrt{k})}$ approximately. Another application of L'Hôpital's rule shows that the limit of the latter is zero.

Turning to (3.A.4c), we have:[17]

$$\begin{aligned} &\frac{f_{h,k}(\bar{s}_k)}{f_{\ell,k}(\bar{s}_k)} \frac{f_{\ell,k}(s_k^\circ)}{f_{h,k}(s_k^\circ)} \\ &\approx \left(\frac{1 - G_h(1 - 1/k)}{1 - G_\ell(1 - 1/k)}\right)\left(\frac{G_\ell(1 + 1/k - 1/\sqrt{k}) - G_\ell(1 - 1/\sqrt{k})}{G_h(1 + 1/k - 1/\sqrt{k}) - G_h(1 - 1/\sqrt{k})}\right). \end{aligned}$$

An application of L'Hôpital's rule shows that the limit of this term equals the limit of

$$\left(\frac{g_h(1)}{g_\ell(1)}\right)\left(\frac{g_\ell(1 + 1/k - 1/\sqrt{k})(1 - \sqrt{k}/2) + g_\ell(1 - 1/\sqrt{k})\sqrt{k}/2}{g_h(1 + 1/k - 1/\sqrt{k})(1 - \sqrt{k}/2) + g_h(1 - 1/\sqrt{k})\sqrt{k}/2}\right).$$

Since g_h and g_ℓ are continuous, this term converges to one, as desired. □

Lemma 3.9. *If s and s′ are signals in S satisfying $f_h(s) \ge f_\ell(s)$ and $s' > s$, then*

$$\frac{1 - \lambda\bar{F}_h(s')}{1 - \lambda\bar{F}_\ell(s')} > \frac{1 - \lambda\bar{F}_h(s)}{1 - \lambda\bar{F}_\ell(s)}$$

holds.

17. When $\{x_k\}_{k\in\mathbb{N}}$ and $\{y_k\}_{k\in\mathbb{N}}$ are sequences of real numbers, we write $x_k \approx y_k$ to mean $\lim_{k\to\infty} x_k/y_k = 1$.

Proof of Lemma 3.9. It suffices to verify this for the case where s' is the signal directly above s. In that case, we have $\underline{F}_\omega(s) + f_\omega(s) = \underline{F}_\omega(s')$. Standard algebraic manipulations show

$$\begin{aligned}&\frac{1-\lambda\bar{F}_h(s')}{1-\lambda\bar{F}_\ell(s')} - \frac{1-\lambda\bar{F}_h(s)}{1-\lambda\bar{F}_\ell(s)}\\ =&\frac{\lambda\left(f_h(s)(1-\lambda\bar{F}_\ell(s)) - f_\ell(s)1-\lambda\bar{F}_h(s)\right)}{\left(1-\lambda\bar{F}_\ell(s')\right)\left(1-\lambda\bar{F}_\ell(s)\right)}.\end{aligned}$$

The (MLRP) implies $(1-\lambda\bar{F}_\ell(s)) \geq (1-\lambda\bar{F}_h(s)$, and we have $f_h(s) \geq f_\ell(s)$ by assumption. □

3.A.3.2 Proof of Proposition 3.3

Proof of Proposition 3.3. We will prove that, for large enough k, there does not exist an equilibrium in which the seller's expected utility equals $\hat{v}_0$. This implies that her utility is bounded away from $\hat{v}_0$ across all equilibria. For, otherwise, compactness of $\Delta(S^\infty) \times \Pi$ lets us extract a convergent subsequence of equilibria along which her utility converges to $\hat{v}_0$; the limit of this subsequence will be an equilibrium in which her utility equals $\hat{v}_0$, and we have a contradiction.

For all k, let $\bar{s}_k = 1 - 1/k$, and let s_k° be as in the conclusion of Lemma 3.8. In what follows, we will suppress the dependence of k from the notation by writing $(S, f_h, f_\ell, s^\circ, \bar{s})$ instead of $(S_k, f_{h,k}, f_{\ell,k}, s_k^\circ, \bar{s}_k)$. No confusion should arise.

In light of Lemma 3.2, we can show that there is no equilibrium where the seller's expected utility equals $\hat{v}_0$ by showing that the pure strategy $\bar{\sigma}$ is not an equilibrium. Towards a contradiction, suppose $\bar{\sigma}$ is supported in equilibrium by some beliefs π of the buyers. We will argue that, for all except finitely-many k, the following strategy constitutes a profitable deviation from $\bar{\sigma}$ for the seller: In the first period, the seller recommends s°; in all later periods i, the seller recommends $\bar{s}$. Let σ denote this sequence of recommendations.

In equilibrium, the first buyer's beliefs do not depend on the seller's action. Thus the prices induced by the deviation are $\hat{v}(s^\circ, \pi_0) = \mathbb{E}[v|\tilde{s} = s^\circ]$ in period 1 and $\hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \bar{\sigma}))$ for all $i \geq 2$. To economize on notation, let

$$x_\omega^\circ \;=\; 1-\lambda\bar{F}_\omega(\bar{s}) \qquad \text{and} \qquad \bar{x}_\omega \;=\; 1-\lambda\bar{F}_\omega(\bar{s}).$$

Thus x_ω° and $\bar{x}_\omega$, respectively, denote the probabilities of not trading after recommending cutoffs s° and $\bar{s}$, respectively, within a given period.

We can now write the seller's utility from the deviation to σ as

$$\begin{aligned}&\mathbb{E}[v|\tilde{s} \;=\; s^\circ] \sum_{\omega\in\{\ell,h\}} \lambda\alpha_{\omega,0}\bar{F}_\omega(\bar{s})\\ &+\sum_{i=2}^{\infty}\left(\hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \bar{\sigma})) \sum_{\omega\in\{\ell,h\}} \lambda\alpha_{\omega,0}\bar{F}_\omega(\bar{s})x_\omega^\circ \bar{x}_\omega^{i-2}\right).\end{aligned} \tag{3.A.5}$$

We complete the proof by arguing that, for all but finitely-many k, the term in the previous expression is strictly greater than $\hat{v}_0$.

Consider the following equality for the expected utility from the deviation (the first expression is simply a restatement of the expected utility from the deviation; the equality adds a zero):

$$\begin{aligned}
&\hat{v}(s^\circ, \pi_0) \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} \bar{F}_\omega(s^\circ) \\
&+ \sum_{i=2}^{\infty} \hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \bar{\sigma})) \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} f_\omega(\bar{s}) x^\circ_\omega \bar{x}^{i-2}_\omega \\
=& \left(\mathbb{E}[v|\tilde{s} = s^\circ] - \mathbb{E}[v|\tilde{s} \geq s^\circ]\right) \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} \bar{F}_\omega(s^\circ) && \text{(3.A.6)} \\
&+ \sum_{i=2}^{\infty} \Bigg(\left(\hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \bar{\sigma})) - \hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \sigma))\right) \\
&\qquad \times \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} f_\omega(\bar{s}) x^\circ_\omega \bar{x}^{i-2}_\omega \Bigg) && \text{(3.A.7)} \\
&+ \mathbb{E}[v|\tilde{s} \geq s^\circ] \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} \bar{F}_\omega(s^\circ) && \text{(3.A.8)} \\
&+ \sum_{i=2}^{\infty} \hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \sigma)) \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} f_\omega(\bar{s}) x^\circ_\omega \bar{x}^{i-2}_\omega && \text{(3.A.9)}
\end{aligned}$$

Iterated expectations show that the sum of (3.A.8) and (3.A.9) equals the prior value $\hat{v}_0$.[18] Thus, to show that utility from the deviation is strictly greater than $\hat{v}_0$, it suffices to show that the sum of (3.A.6) and (3.A.7) is strictly positive.

Several lines of algebra establish the following identities:

$$\begin{aligned}
&\mathbb{E}[v|\tilde{s} = s^\circ] - \mathbb{E}[v|\tilde{s} \geq s^\circ] \\
=& \frac{(v_h - v_\ell)\alpha_h \alpha_\ell \left(f_h(s^\circ)\bar{F}_\ell(s^\circ) - f_\ell(s^\circ)\bar{F}_h(s^\circ)\right)}{\left(\sum_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(s^\circ)\right)\left(\sum_{\omega \in \{\ell,h\}} \alpha_{\omega,0} \bar{F}_\omega(s^\circ)\right)},
\end{aligned}$$

and

$$\begin{aligned}
&\hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \bar{\sigma})) - \hat{v}(\bar{s}, \hat{\pi}_i(\bar{s}, \sigma)) \\
=& \frac{(v_h - v_\ell)\alpha_h \alpha_\ell f_h(\bar{s}) f_\ell(\bar{s}) \bar{x}^{i-2}_h \bar{x}^{i-2}_\ell \left(\bar{x}_h x^\circ_\ell - \bar{x}_\ell x^\circ_h\right)}{\left(\sum_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s}) \bar{x}^{i-1}_\omega\right)\left(\sum_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s}) x^\circ_\omega \bar{x}^{i-2}_\omega\right)}.
\end{aligned}$$

18. For each $i \geq 2$, the summand in (3.A.9) is the probability that trade happens in period i under the sequence σ. multiplied by the posterior value conditional on said event. In (3.A.8), we note that $\mathbb{E}[v|\tilde{s} \geq s^\circ] \sum_{\omega \in \{\ell,h\}} \lambda \alpha_{\omega,0} \bar{F}_\omega(s^\circ)$ is precisely that trade happens in period 1 under σ multiplied by the posterior value on that event.

Notice that the positive term $(v_h - v_\ell)\alpha_h\alpha_\ell$ appears in both of the previous two identities. For the purposes of evaluating the sign of the sum of (3.A.6) and (3.A.7), we may ignore this term. If we now plug the previous two identities back into (3.A.6) and (3.A.7), it follows that we must verify that the following sum is strictly positive sufficiently far enough along the sequence of signal structures:

$$\frac{\left(f_h(s^\circ)\bar{F}_\ell(s^\circ) - f_\ell(s^\circ)\bar{F}_h(s^\circ)\right)}{\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0} f_\omega(s^\circ)} \tag{3.A.10}$$

$$+\sum_{i=2}^{\infty} \frac{f_h(\bar{s})f_\ell(\bar{s})\bar{x}_h^{i-2}\bar{x}_\ell^{i-2}\left(\bar{x}_h x_\ell^\circ - \bar{x}_\ell x_h^\circ\right)}{\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})\bar{x}_\omega^{i-1}}. \tag{3.A.11}$$

For convenience, let us restate the implications of Lemma 3.8 (the dependence on k being suppressed in the notation).

$$\infty > \lim_{k\to\infty} \frac{f_h(\bar{s})}{f_\ell(\bar{s})} > 1, \tag{3.A.12a}$$

$$\lim_{k\to\infty} \frac{f_h(\bar{s})}{1 - \underline{F}_\ell(s^\circ)} = 0, \tag{3.A.12b}$$

$$\lim_{k\to\infty} \frac{f_h(\bar{s})}{f_\ell(\bar{s})} \frac{f_\ell(s^\circ)}{f_h(s^\circ)} = 1. \tag{3.A.12c}$$

We continue by establishing a lower bound on the term in (3.A.11). Consider the difference

$$\begin{aligned} &\bar{x}_h x_\ell^\circ - \bar{x}_\ell x_h^\circ \\ =&(1 - \lambda f_h(\bar{s}))(1 - \lambda\bar{F}_\ell(s^\circ)) - (1 - \lambda f_\ell(\bar{s}))(1 - \lambda\bar{F}_h(s^\circ). \end{aligned} \tag{3.A.13}$$

We claim that this difference is strictly positive for all large k. We know from (3.A.12a) that $f_h(\bar{s})/f_\ell(\bar{s})$ is strictly greater than one and eventually bounded away from one. Further, $\frac{f_h(\bar{s})}{f_\ell(\bar{s})}\frac{f_\ell(s^\circ)}{f_h(s^\circ)}$ approaches one. Thus $f_h(s^\circ)/f_\ell(s^\circ)$ must be strictly larger than one for all sufficiently large k. It now follows from Lemma 3.9 that (3.A.13) is strictly positive for such k.

Next, consider the ratio

$$\frac{\bar{x}_\ell^{i-2}}{\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})\bar{x}_\omega^{i-1}}. \tag{3.A.14}$$

Recall the definition $x_\omega(s) = 1 - \lambda(1 - \underline{F}_\omega(s))$. The (MLRP) implies $\bar{x}_h^{i-1}/\bar{x}_\ell^{i-2} \le 1$, and hence the following is a lower bound on (3.A.14):

$$\frac{\bar{x}_\ell^{i-2}}{\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})\bar{x}_\omega^{i-1}} \ge \frac{1}{\sum_{\omega\in\{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})}. \tag{3.A.15}$$

The fact that the term in (3.A.13) is strictly positive and the inequality in (3.A.15) together imply that the following is a lower bound on (3.A.11):

$$\begin{aligned}
&\sum_{i=2}^{\infty} \frac{f_h(\bar{s}) f_\ell(\bar{s}) \bar{x}_h^{i-2} \bar{x}_\ell^{i-2} \left(\bar{x}_h x_\ell^\circ - \bar{x}_\ell x_h^\circ\right)}{\sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s}) \bar{x}_\omega^{i-1}} \\
\geq & \frac{f_h(\bar{s}) f_\ell(\bar{s})}{\sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})} \left(\bar{x}_h x_\ell^\circ - \bar{x}_\ell x_h^\circ\right) \sum_{i=2}^{\infty} \bar{x}_h^{i-2} \\
= & \frac{f_h(\bar{s}) f_\ell(\bar{s})}{\sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})} \left(\bar{x}_h x_\ell^\circ - \bar{x}_\ell x_h^\circ\right) \frac{1}{1-\bar{x}_h}.
\end{aligned}$$

If we plug back in the definition $x_\omega(s) = 1 - \lambda(1 - \underline{F}_\omega(s)) = 1 - \lambda \bar{F}_\omega(s)$, we obtain

$$f_\ell(\bar{s}) \frac{(1-\lambda f_h(\bar{s}))(1-\lambda \bar{F}_\ell(s^\circ)) - (1-\lambda f_\ell(\bar{s}))(1-\lambda \bar{F}_h(s^\circ))}{\lambda \sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})}. \tag{3.A.16}$$

To summarize: We may complete the proof by verifying that the sum of (3.A.10) and (3.A.16) is strictly positive for all sufficiently large k. This sum reads

$$\frac{\left(f_h(s^\circ)\bar{F}_\ell(s^\circ) - f_\ell(s^\circ)\bar{F}_h(s^\circ)\right)}{\sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(s^\circ)} \tag{3.A.17}$$

$$+ f_\ell(\bar{s}) \frac{(1-\lambda f_h(\bar{s}))(1-\lambda \bar{F}_\ell(s^\circ)) - (1-\lambda f_\ell(\bar{s}))(1-\lambda \bar{F}_h(s^\circ))}{\lambda \sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})}. \tag{3.A.18}$$

It is useful to rearrange the sum of (3.A.17) and (3.A.18) before proceeding.

Dividing the sum of (3.A.17) and (3.A.18) by

$$\frac{f_\ell(\bar{s})}{\bar{F}_\ell(s^\circ) \sum\limits_{\omega \in \{\ell,h\}} \alpha_{\omega,0} f_\omega(\bar{s})}$$

leaves its sign unchanged. Rearranging the resulting terms further via standard algebraic manipulations, we find that the sign of the sum of (3.A.17) and (3.A.18) is the sign of

$$
\begin{aligned}
&\left(\frac{f_h(s^\circ)}{f_\ell(\bar{s})}\frac{\sum\limits_{\omega\in\{\ell,h\}}\alpha_{\omega,0}f_\omega(\bar{s})}{\sum\limits_{\omega\in\{\ell,h\}}\alpha_{\omega,0}f_\omega(s^\circ)}-1\right)\\
&-\frac{\bar{F}_h(s^\circ)}{\bar{F}_\ell(s^\circ)}\left(\frac{f_\ell(s^\circ)}{f_\ell(\bar{s})}\frac{\sum\limits_{\omega\in\{\ell,h\}}\alpha_{\omega,0}f_\omega(\bar{s})}{\sum\limits_{\omega\in\{\ell,h\}}\alpha_{\omega,0}f_\omega(s^\circ)}-1\right)\\
&+\frac{f_\ell(\bar{s})(1-\lambda\bar{F}_h(s^\circ))-f_h(\bar{s})(1-\lambda\bar{F}_\ell(s^\circ))}{\bar{F}_\ell(s^\circ)}\\
&=\left(\pi_0\left(\frac{f_h(\bar{s})}{f_\ell(\bar{s})}-1\right)+1-\frac{f_\ell(s^\circ)}{f_h(s^\circ)}\right)\left(\frac{f_\ell(s^\circ)}{f_h(s^\circ)}+\pi_0\right)^{-1} && (3.A.19)\\
&-\pi_0\frac{\bar{F}_h(s^\circ)}{\bar{F}_\ell(s^\circ)}\frac{f_h(s^\circ)}{f_\ell(s^\circ)}\left(\frac{f_h(\bar{s})}{f_\ell(\bar{s})}\frac{f_\ell(s^\circ)}{f_h(s^\circ)}-1\right)\left(1+\pi_0\frac{f_h(s^\circ)}{f_\ell(s^\circ)}\right)^{-1} && (3.A.20)\\
&+\frac{f_\ell(\bar{s})(1-\lambda\bar{F}_h(s^\circ))-f_h(\bar{s})(1-\lambda\bar{F}_\ell(s^\circ))}{\bar{F}_\ell(s^\circ)}. && (3.A.21)
\end{aligned}
$$

We complete the proof by arguing that, along the sequence of signal structures, the term in (3.A.19) is positive (far enough along the sequence) and bounded away from 0, whereas the sum of (3.A.20) and (3.A.21) admits a lower bound that converges to 0.

Beginning with (3.A.19) we infer from (3.A.12a) and (3.A.12b) that $\frac{f_h(\bar{s})}{f_\ell(\bar{s})}-1$ and $1-\frac{f_\ell(s^\circ)}{f_h(s^\circ)}$ are eventually and bounded away from 0. We also know from (3.A.12a) that $\left(\frac{f_\ell(s^\circ)}{f_h(s^\circ)}+\pi_0\right)^{-1}$ is bounded. Hence (3.A.19) is eventually positive and bounded away from 0.

Turning to (3.A.19), we infer from (3.A.21) that $\left(\frac{f_h(\bar{s})}{f_\ell(\bar{s})}\frac{f_\ell(s^\circ)}{f_h(s^\circ)}-1\right)$ converges to 0. The ratio $\frac{\bar{F}_h(s^\circ)}{\bar{F}_\ell(s^\circ)}$ is bounded (it converges to the ratio of the densities at 1.) Simultaneously, we know from (3.A.12a) and (3.A.12c) that all others terms in (3.A.19) are bounded along the sequence. Thus (3.A.20) converges to 0.

Lastly, turning to (3.A.21), we have the following lower bound on (3.A.21):

$$\frac{f_\ell(\bar{s})(1-\lambda\bar{F}_h(s^\circ))-f_h(\bar{s})(1-\lambda\bar{F}_\ell(s^\circ))}{\bar{F}_\ell(s^\circ)} \geq -\frac{f_h(\bar{s})(1-\lambda\bar{F}_\ell(s^\circ))}{\bar{F}_\ell(s^\circ)}.$$

We conclude from (3.A.12c) that this lower bound converges to 0. □

3.A.4 Surplus extraction with binary signals

Proof of Proposition 3.4. We proceed along a number of claims.

Claim 3.10. *If $\mu'\in\Delta(S^\infty)$ is a fully mixed mixed strategy, then for all i and s we have $\hat{\pi}_i(s,\mu')\geq\hat{\pi}_i(\bar{s},\bar{\sigma})$*

Proof of Claim 3.10. Given an integer i and a pure strategy σ, let $N(i,\sigma) = |\{j \in \{1,\ldots,i-1\} : \sigma_j = \bar{s}\}|$. That is, $N(i,\sigma)$ is the number of times σ plays $\bar{s}$ in rounds 1 to $i-1$.

Since signals are binary, we may write the posterior $\hat{\pi}_i(s,\mu')$ as follows:

$$\begin{aligned}\hat{\pi}_i(s,\mu') &= \pi_0 \frac{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\prod_{j=1}^{i-1}(1-\lambda(1-\underline{F}_h(\sigma_j)))}{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\prod_{j=1}^{i-1}(1-\lambda(1-\underline{F}_\ell(\sigma_j)))}\\
&= \pi_0 \frac{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\,(1-\lambda)^{(i-1)-N(i,\sigma)}\left(1-\lambda f_h(\bar{s})\right)^{N(i,\sigma)}}{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\,(1-\lambda)^{(i-1)-N(i,\sigma)}\left(1-\lambda f_\ell(\bar{s})\right)^{N(i,\sigma)}}\\
&= \pi_0 \frac{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\left(\frac{1-\lambda f_h(\bar{s})}{1-\lambda}\right)^{N(i,\sigma)}}{\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\left(\frac{1-\lambda f_\ell(\bar{s})}{1-\lambda}\right)^{N(i,\sigma)}}.\end{aligned}$$

The belief $\hat{\pi}_i(\bar{s},\bar{\sigma})$ is given by

$$\hat{\pi}_i(\bar{s},\bar{\sigma}) = \pi_0\left(\frac{1-\lambda f_h(\bar{s})}{1-\lambda f_\ell(\bar{s})}\right)^{i-1}.$$

Hence the sign of difference $\hat{\pi}_i(s,\mu') - \hat{\pi}_i(\bar{s},\bar{\sigma})$ is the sign of

$$\begin{aligned}&\sum_{\sigma\in S^n} \mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)\Bigg(\left(1-\lambda f_\ell(\bar{s})\right)^{i-1}\left(\frac{1-\lambda f_h(\bar{s})}{1-\lambda}\right)^{N(i,\sigma)}\\
&\qquad\qquad -\left(1-\lambda f_h(\bar{s})\right)^{i-1}\left(\frac{1-\lambda f_\ell(\bar{s})}{1-\lambda}\right)^{N(i,\sigma)}\Bigg)\\
&= \sum_{\sigma\in S^n}\Bigg(\frac{\mathbf{1}_{(\sigma_i=s)}\mu'(\sigma)}{(1-\lambda)^{N(i,\sigma)}}\left(\frac{(1-\lambda f_h(\bar{s}))(1-\lambda f_\ell(\bar{s}))}{1-\lambda}\right)^{N(i,\sigma)}\\
&\qquad\qquad \times\left(\left(1-\lambda f_\ell(\bar{s})\right)^{i-1-N(i,\sigma)} - \left(1-\lambda f_h(\bar{s})\right)^{i-1-N(i,\sigma)}\right)\Bigg).\end{aligned}$$

This sum is weakly positive since $f_\ell(\bar{s}) < f_h(\bar{s})$ holds and since $N(i,\sigma)$ is no greater than $i-1$. □

Claim 3.11. *If (μ,π) is a sequential equilibrium of $\Gamma(\infty,0)$, then the seller's equilibrium expected utility is $\mathbb{E}[v|\pi_0]$, and μ is the pure strategy $\bar{\sigma}$*

Proof of Claim 3.11. Since (μ,π) is a sequential equilibrium, the belief π is the pointwise limit of a sequence of Bayesian posteriors derived from fully-mixed strategies. Thus Claim 3.10 implies that, for arbitrary $i \in \mathbb{N}$ and $s \in S$, the belief $\pi_i(s)$ is bounded below by $\hat{\pi}_i(\bar{s},\bar{\sigma})$. The seller's utility is pointwise-increasing in the beliefs. Hence, given beliefs π, and the seller's utility from the pure strategy $\bar{\sigma}$ is at least $V(\bar{\sigma}, \hat{\pi}(\cdot, \bar{\sigma}))$. Since $\hat{\pi}_i(\bar{s}, \bar{\sigma})$ is the posterior induced by $\bar{\sigma}$, the reasoning of Lemma 3.2 via iterated expectations shows that $V(\bar{\sigma}, \hat{\pi}(\cdot, \bar{\sigma}))$ equals $\hat{v}_0$. Thus $\hat{v}_0$ is a lower bound on the seller's equilibrium utility. We know from Lemma 3.2 that $\hat{v}_0$ is also an upper bound on the seller's equilibrium utility. Hence another application of Lemma 3.2 shows that the seller's utility is $\hat{v}_0$ and that her strategy is $\bar{\sigma}$. □

□

Appendix 3.B Unobservable time-on-the-market

3.B.1 Definitions and notation

In this section we derive expressions for the seller's expected utility and buyers' posterior beliefs in the game of Section 3.3.

Let $n \in \mathbb{N}$ and let $c \in [0, \lambda v_\ell]$. A mixed strategy of the seller is an element μ_n of $\Delta(S^n)$. Buyers' beliefs are represented by a function $\pi_n^\emptyset \colon S \to [0, \pi_0]$. (In the main text, we introduced beliefs as a function mapping to $[0, \infty]$, but, as in Appendix 3.A.1, it is without loss to focus on beliefs in $[0, \pi_0]$.)

3.B.1.1 The seller's expected utility

Given $\pi_n^\emptyset$, the seller's expected utility from a mixed strategy μ_n is

$$V^\emptyset(\mu_n, \pi_n^\emptyset, n, c) = \sum_{\sigma \in S^n} \mu_n(\sigma) \sum_{i=1}^{n} \sum_{\omega \in \{\ell, h\}} \Bigg(\alpha_\omega \prod_{j=1}^{i-1} \big(1 - \lambda \bar{F}_\omega(\sigma_j)\big) \times \big(\lambda \bar{F}_\omega(\sigma_i) \hat{v}(\sigma_i, \pi^\emptyset(\sigma_i)) - c\big) \Bigg). \tag{3.B.1}$$

3.B.1.2 Buyers' inference

Given a mixed strategy μ_n, let $S_n(\mu_n)$ denote the subset of signals that μ_n plays with non-zero probability in at least one of the n periods. That is, s is in $S_n(\mu_n)$ if and only if

$$\sum_{i=1}^{n} \sum_{\sigma \in S^n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) > 0. \tag{3.B.2}$$

Given μ_n and $s \in S_n(\mu_n)$, the Bayesian posterior conditional on arriving to the market on being recommended s is well-defined. We denote it by $\hat{\pi}_n^\emptyset(s, \mu_n)$. It is given by

$$\hat{\pi}_n^\emptyset(s, \mu_n) = \pi_0 \frac{\sum_{\sigma \in S^n} \sum_{i=1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \big(1 - \lambda \bar{F}_\omega(\sigma_j)\big)}{\sum_{\sigma \in S^n} \sum_{i=1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \big(1 - \lambda \bar{F}_\omega(\sigma_j)\big)}. \tag{3.B.3}$$

A mixed strategy μ_n is *fully mixed* if $S_n(\mu_n) = S$; that is, if each cutoff is recommended with non-zero probability in at least one period.

3.B.2 Auxiliary Results

This part of the appendix presents some auxiliary results.

Lemma 3.12. *Let $n \in \mathbb{N}$, $c \in [0, \lambda v_\ell]$, and $\mu_n \in \Delta(S^n)$. Let $\pi_n^\emptyset : S \to [0,1]$ be a function that agrees with $\hat{\pi}_n^\emptyset(\cdot, \mu_n)$ at all s in $S_n(\mu_n)$. Then we have*

$$V^\emptyset(\mu_n, \pi_n^\emptyset, n, c) \;\leq\; \hat{v}_0.$$

The proof is analogous to that of Lemma 3.2 and is omitted.

The next result is chiefly used in the upcoming proof of Proposition 3.6; the reader may prefer to skip the result for now returning to it as needed. Consider an auxiliary fixed-point problem in which, for some given integer j and signal s^*, the seller is restricted to randomizing over pure strategies in which s^* is played in all of the first j rounds. Formally, given $n \in \mathbb{N}$ and $j \in \mathbb{N} \cup \{0\}$ such that $n-1 \geq j$, let

$$\Sigma_{n,j,s^*} \;=\; \left\{(\sigma_1, \ldots, \sigma_n) \;\in\; S^n \colon \left(\forall_{k' \colon 1 \leq k' \leq j},\, \sigma_{k'} \;=\; s^*\right)\right\}.$$

The set of probability distributions over Σ_{n,j,s^*} is denoted by $\Delta(\Sigma_{n,j,s^*})$. As a convention, for $j = 0$, the set Σ_{n,j,s^*} means the set S^n.

Lemma 3.13. *Let $c \in [0, v_\ell]$. Let $n \in \mathbb{N}$ and $j \in \mathbb{N} \cup \{0\}$ be such that $n-1 \geq j$. Let $s^* \in S$. There exists a sequence $\{\mu_k\}_{k \in \mathbb{N}}$ in $\Delta(\Sigma_{n,j,s^*})$, a strategy $\mu_{n,j}$ in $\Delta(\Sigma_{n,j,s^*})$, and a belief $\pi_{n,j}^\emptyset$ satisfying all of the following:*

(1) We have

$$\mu_{n,j} \;\in\; \operatorname*{arg\,max}_{\mu' \in \Delta(\Sigma_{n,j,s^*})} V^\emptyset(\mu', \pi_{n,j}^\emptyset, n, c). \tag{3.B.4}$$

(2) For all k, the strategy μ_k is fully mixed.

(3) The sequence $\{\mu_k\}_{k \in \mathbb{N}}$ converges to $\mu_{n,j}$ as $k \to \infty$.

(4) The sequence of induced beliefs $\{\hat{\pi}_n^\emptyset(\cdot, \mu_k)\}_k$ converges to $\pi_{n,j}^\emptyset$ as $k \to \infty$.

The proof proceeds via routine arguments and is omitted.

As an immediate corollary, we find that $\Gamma^\emptyset(n, c)$ admits some sequential equilibrium. In the main text, this was stated as Lemma 3.5.

Proof of Lemma 3.5. Invoke Lemma 3.13 with $j = 0$. □

The next auxiliary lemma will be useful for the upcoming proof of Proposition 3.6; the reader may prefer to skip the result for now returning to it as needed. It characterizes the beliefs which are induced by a strategy in $\Delta(\Sigma_{n,j,s^*})$ for large n and j. Verbally, all on-path cutoffs different from s^* lead to a belief that the state is ℓ with overwhelming probability. Conversely, the belief at s^* is approximately $\pi_0 \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)}$.

Lemma 3.14. *Let $s^* \in S \setminus \{\underline{s}\}$. Let $\{j_n\}_{n\in\mathbb{N}}$ be a sequence of integers. For all n, let μ_n be a mixed strategy in $\Delta(\Sigma_{n,j_n,s^*})$. If the sequence $(j_n)_{n\in\mathbb{N}}$ diverges to $+\infty$, then for every $\varepsilon > 0$ there exists $n_\varepsilon \in \mathbb{N}$ such that for all n greater than n_ε all of the following are true:*

(1) If s is in $S_n(\mu_n) \setminus \{s^*\}$, then $\hat{\pi}_n^{\emptyset}(s^*, \mu_n) < \varepsilon$ holds.
(2) We have

$$\left| \hat{\pi}_n^{\emptyset}(s^*, \mu_n) - \pi_0 \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} \right| < \varepsilon. \tag{3.B.5}$$

Proof of Lemma 3.14. Let $\varepsilon > 0$. Turning to the first claim, let n be arbitrary and consider a signal s in $S_n(\mu_n) \setminus \{s^*\}$. The posterior $\hat{\pi}_n^{\emptyset}(s, \mu_n)$ is defined to be

$$\hat{\pi}_n^{\emptyset}(s, \mu_n) = \frac{\sum\limits_{\sigma\in S^n} \sum\limits_{i=1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_h(\sigma_j)\right)}{\sum\limits_{\sigma\in S^n} \sum\limits_{i=1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_\ell(\sigma_j)\right)}.$$

By definition of $\Delta(\Sigma_{n,j_n,s^*})$, the distribution μ_n assigns positive probability to a pure strategy σ only if σ is in Σ_{n,j_n,s^*}. Accordingly, conditional on seeing a signal different from s^*, a buyer can be sure that at least j_n rounds have passed in which s^* was not accepted. This implies the following identity for arbitrary ω:

$$\begin{aligned}
&\sum_{\sigma\in S^n} \sum_{i=1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_\omega(\sigma_j)\right) \\
&= \left(1 - \lambda \bar{F}_\omega(s^*)\right)^{j_n} \sum_{\sigma\in \Sigma_{n,j_n,s^*}} \sum_{i=j_n+1}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=j_n+1}^{i-1} \left(1 - \lambda \bar{F}_\omega(\sigma_j)\right).
\end{aligned}$$

Hence the posterior belief $\hat{\pi}_n^{\emptyset}(s, \mu_n)$ reads

$$\hat{\pi}_n^{\emptyset}(s, \mu_n) = \left(\frac{1 - \lambda \bar{F}_h(s^*)}{1 - \lambda \bar{F}_\ell(s^*)}\right)^{j_n - 1} \frac{\sum\limits_{\sigma\in \Sigma_{n,j_n,s^*}} \sum\limits_{i=j_n}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_h(\sigma_j)\right)}{\sum\limits_{\sigma\in \Sigma_{n,j_n,s^*}} \sum\limits_{i=j_n}^{n} \mathbf{1}_{(\sigma_i = s)} \mu_n(\sigma) \prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_\ell(\sigma_j)\right)}.$$

The (MLRP) implies that the second fraction in this expression is less than 1. Moreover, since s^* is not $\underline{s}$, we infer from the (MLRP) that $1 - \lambda \bar{F}_h(s^*) < 1 - \lambda \bar{F}_\ell(s^*)$ holds. Thus there is some integer j'_ε satisfying

$$j_n \geq j'_\varepsilon \quad \Rightarrow \quad \left(\frac{1 - \lambda \bar{F}_h(s^*)}{1 - \lambda \bar{F}_\ell(s^*)}\right)^{j_n - 1} < \varepsilon.$$

In particular, for such j_n above j'_ε, the belief $\hat{\pi}_n^{\emptyset}(s, \mu_n)$ is less than ε for all $s \in S_n(\mu_n)$. Keeping this in mind, let us turn to the second part of the claim.

Consider the probability that a buyer assigns to following joint event: He arrives to the market when the object has not yet been traded and is then offered a signal of s^*. Conditional on the state being ω, we denote this probability by $q_{\omega,n}$; it is given by

$$q_{\omega,n} = \frac{1}{n}\sum_{\sigma\in S^n}\sum_{i=1}^{n}\mathbf{1}_{(\sigma_i=s^*)}\mu_n(\sigma)\prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_j)\big).$$

Using that μ_n is in $\Delta(\Sigma_{n,j_n,s^*})$, we find that $q_{\omega,n}$ equals

$$\frac{1}{n}\Bigg(\sum_{i=1}^{j_n}\big(1-\lambda\bar{F}_\omega(s^*)\big)^{i-1} + \big(1-\lambda\bar{F}_\omega(s^*)\big)^{j_n}\sum_{\sigma\in\Sigma_{n,j_n,s^*}}\sum_{i=j_n+1}^{n}\mu_n(\sigma)\mathbf{1}_{(\sigma_i=s^*)}\prod_{j=j_n+1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_j)\big)\Bigg)$$

A moment's thought reveals that the following are lower and upper bounds, respectively, on $q_{\omega,n}$:

$$\begin{aligned}&\frac{1}{n}\sum_{i=1}^{j_n}\big(1-\lambda\bar{F}_\omega(s^*)\big)^{i-1}\\ \le\, & q_{\omega,n}\\ \le\, & \frac{1}{n}\left(\sum_{i=1}^{j_n}\big(1-\lambda\bar{F}_\omega(s^*)\big)^{i-1}\right)+\big(1-\lambda\bar{F}_\omega(s^*)\big)^{j_n}\sum_{i=j_n+1}^{n}(1-\lambda\bar{F}_\omega(\bar{s}))^{i-(j_n+1)}.\end{aligned}$$

Recall that $j_n\to\infty$ as $n\to\infty$. Hence we have $nq_{\omega,n}\to 1/(\lambda\bar{F}_\omega(s^*)$ as $n\to\infty$. The posterior belief $\hat{\pi}_n^{\emptyset}(s^*,\mu_n)$ is equal to the ratio $\pi_0 q_{h,n}/q_{\ell,n}$. Hence, there is an integer j''_ε such that $\hat{\pi}_n^{\emptyset}(s^*,\mu_n)$ is within ε of $\pi_0\bar{F}_\ell(s^*)/\bar{F}_h(s^*)$ if j_n is greater than j''_ε .

Let $j_\varepsilon=\max(j'_\varepsilon,j''_\varepsilon)$. The preceding arguments show that all desired inequalities hold for μ_n if j_n is above j_ε. Recalling that $j_n\to\infty$ as $n\to\infty$, the claim follows. □

3.B.3 Signaling calendar time

This part of the appendix is devoted to the proof of Proposition 3.6. Before delving into the details, let us sketch the idea. Let $s^*\in S\setminus\{\underline{s}\}$. We begin by defining a sequence $(j_n)_{n\in\mathbb{N}}$ of integers. Think of this as a sequence that diverges to ∞, but not too rapidly. In the game with n buyers, we then consider a restricted notion of equilibrium in which the seller is required to play s^* in the first j_n rounds (including in deviations). Lemma 3.13 from Appendix 3.B.2 shows that such an "equilibrium" exists. As long as j_n diverges to ∞, Lemma 3.14 from Appendix 3.B.2 then implies that all signals different from s^* induce buyers to update their beliefs such that their willingness to pay is approximately v_ℓ; this step requires that s^* be different from

$\underline{s}$. Moreover, their willingness to pay after s^* is bounded away from v_ℓ; it is approximately the expression for the seller's limit utility as given in (3.5). We can then show that, far enough along the sequence, the seller will indeed find it optimal to only play s^* in the first, say, i_n rounds. To complete the proof, it is therefore sufficient to check that i_n is eventually larger than j_n, i.e. that the earlier constraint on the seller's strategy is eventually non-binding. For this final step, we require that j_n not diverge too quickly. Intuitively, the seller may have a motive to save on search costs even if this entails trading at an undesirable price. Thus the speed at which costs c_n vanish needs to be taken into account when letting j_n diverge.

Proof of Proposition 3.6. We shall first define the sequence of integers $\{j_n\}_{n\in\mathbb{N}}$. Our candidates for $(\mu_n, \pi_n^\emptyset)$ will then be derived from Lemma 3.13.

For all n let j_n denote the largest integer (or zero) smaller than

$$\min\left(\frac{1}{2}\frac{\ln v_\ell - \ln c_n}{-\ln(1-\lambda\bar{F}_h(s^*))}, \frac{n}{2}\right). \tag{3.B.6}$$

(If $c_n = 0$, we understand the minimum to equal $n/2$.) The motivation for this obscure choice of j_n will reveal itself in the final step of the proof. For the moment, we only note that $n-1 \geq j_n$ holds, and that j_n and $n-j_n$ both go to infinity as n goes to infinity.

For arbitrary n, we may appeal to Lemma 3.13 with j_n in the role of j to assert the following:

Claim 3.15. *For all n, there exists a strategy μ_n in Σ_{n,j_n,s^*}, a belief $\pi_n^\emptyset$, and a sequence $\{\mu_{n,k}\}_{k\in\mathbb{N}}$ in $\Delta(\Sigma_{n,j_n,s^*})$ such that all of the following are true:*

(1) We have

$$\mu_n \in \underset{\mu'\in\Delta(\Sigma_{n,j_n,s^*})}{\arg\max}\; V^\emptyset(\mu', \pi^\emptyset_{n,j_n}, n, c). \tag{3.B.7}$$

(2) For all k, the strategy $\mu_{n,k}$ is completely mixed, i.e. the sets $S_n(\mu_{n,k})$ and S are equal.

(3) The sequence $(\mu_{n,k})_{k\in\mathbb{N}}$ converges to μ_n as $k \to \infty$.

(4) The sequence of induced beliefs $(\hat{\pi}_n^\emptyset(\cdot, \mu_{n,k}))_k$ converges to $\pi_n^\emptyset$ as $k \to \infty$.

In what follows, we understand that for all n, the tuple $(\mu_n, \pi_n^\emptyset, \{\mu_{n,k}\}_{k\in\mathbb{N}})$ is as in the conclusion of Equation (3.B.7). We will prove that the sequence $\{\mu_n, \pi_n^\emptyset, j_n\}_{n\in\mathbb{N}}$ satisfies all desired properties.

As a first step, we use Lemma 3.14 to characterize the beliefs $\pi_n^\emptyset$. Note that Lemma 3.14 is silent on the beliefs at off-path cutoffs. We will make use of the fact that, by construction, μ_n and $\pi_n^\emptyset$ are well-approximated by $\mu_{n,k}$ and $\hat{\pi}_n^\emptyset(\mu_{n,k})$, respectively. Since $\mu_{n,k}$ is completely mixed, we may then use Lemma 3.14 to characterize the beliefs at all cutoffs.

Claim 3.16. *For all $\varepsilon > 0$ there exists $n_\varepsilon \in \mathbb{N}$ such that for all n greater than n_ε all of the following are true:*

(1) If s is in $S \setminus \{s^*\}$, then $\hat{\pi}_n^\emptyset(s, \mu_n) < \varepsilon$ holds.
(2) We have

$$\left| \hat{\pi}_n^\emptyset(s^*, \mu_n) - \pi_0 \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} \right| < \varepsilon.$$

Proof Claim 3.16. Let $\varepsilon > 0$. Part 4. of Claim 3.15 implies that for all n we may find k_n such that $|\hat{\pi}^\emptyset(s, \mu_{n,k_n}) - \pi_n^\emptyset(s)| < \varepsilon/2$ holds for all $s \in S$. For later reference, note that μ_{n,k_n} is fully mixed.

Consider the sequence $\{\mu_{n,k_n}\}_{n\in\mathbb{N}}$ thus defined. Since j_n diverges to ∞, we may appeal to Lemma 3.14 to find an integer n_ε such that for all n above n_ε all of the following are true:

(1) If s is in $S_n(\mu_{n,k_n}) \setminus \{s^*\}$, then $\hat{\pi}^\emptyset(s, \mu_{n,k_n}) < \varepsilon/2$.
(2) We have

$$\left| \hat{\pi}^\emptyset(s^*, \mu_{n,k_n}) - \pi_0 \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} \right| < \varepsilon/2.$$

Since μ_{n,k_n} is fully mixed, the sets $S_n(\mu_{n,k_n})$ and S are equal. We also recall that the inequality

$$|\hat{\pi}^\emptyset(s, \mu_{n,k_n}) - \pi_n^\emptyset(s)| < \varepsilon/2$$

holds for all $s \in S$. The claim follows from the above inequalities. □

The previous step allows an easy comparison of the prices that the seller can hope to obtain under $\pi_n^\emptyset$. To keep some of more algebraic steps readable, we simplify notation. For all $s \in S$ let

$$u_n(s) \ = \ \hat{v}(s, \pi_n^\emptyset(s))$$

denote the price that the seller would obtain from trading at a cutoff of s when beliefs are $\pi_n^\emptyset$. An immediate implication of Claim 3.16 is that $u_n(s)$ converges to v_ℓ for all s different from s^*. Moreover, we have

$$u_n(s^*) \xrightarrow{n\to\infty} \frac{1}{\pi_0 \frac{f_h(s^*)}{f_\ell(s^*)} \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} + 1} \left(v_h \pi_0 \frac{f_h(s^*)}{f_\ell(s^*)} \frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} + v_\ell \right).$$

In particular, for all sufficiently large values of n and all s different from s^*, we may assert that $u_n(s^*) - u_n(s)$ is positive and bounded away from zero. These inequalities and the limit for $u_n(s^*)$ are the only properties of $\pi_n^\emptyset$ that will be relevant in the remainder of the proof.

In what follows, if σ_n in S^n is some pure strategy for some n, then $\sigma_{n,i}$ means the i'th entry of σ_n, i.e. the seller's action in period i.

Claim 3.17. *There exists $n^* \in \mathbb{N}$ such that for all n greater than n^*, if $\sigma_n$satisfies*

$$\sigma \in \arg\max_{\sigma' \in S^n} V^{\emptyset}(\sigma', \pi_n^{\emptyset}, n, c_n),$$

then the following is true: For all i and i', if $\sigma_{n,i} = s^*$ and $i' < i$, then $\sigma_{n,i'} = s^*$.

In other words, eventually, every pure best response to $\pi_n^{\emptyset}$ will admit a cutoff-structure: If s^* is played, then it is played up to some integer-cutoff, and never afterwards. (This integer-cutoff may be different for each best response. The cutoff may also equal n, in which case the strategy plays s^* in all periods.)

Proof Claim 3.17. Recall that $u_n(s^*) - u_n(s)$ is positive and bounded away from zero for all sufficiently large values of n and all $s \in S \setminus \{s^*\}$. Recall also that c_n converges to zero. Hence we may find an integer n^* such if n is greater than n^* and s is in $S \setminus \{s^*\}$, then for all ω the inequality

$$u_n(s^*) - u_n(s) > \frac{c_n}{\lambda} \frac{\underline{F}_\omega(s^*) - \underline{F}_\omega(s)}{\bar{F}_\omega(s)\bar{F}_\omega(s^*)}$$

holds.

Fix an integer n greater than n^*, and let σ be in $\in \arg\max_{\sigma' \in S^n} V^{\emptyset}(\sigma', \pi_n^{\emptyset}, n, c_n)$. Towards a contradiction, suppose the claim was false. Then there exists an index i satisfying $\sigma_{n,i+1} = s^*$ and $\sigma_{n,i} \neq s^*$. Let $v_{i+2,\omega}$ denote the seller's expected payoff from period $i+2$ onwards under σ, conditional on state ω. Conditional on state ω, her expected payoff from period i onwards under σ is thus given by

$$\begin{aligned} & u_n(\sigma_{n,i})\lambda\bar{F}_\omega(\sigma_{n,i}) \\ & + \big(u_n(s^*)\lambda\bar{F}_\omega(s^*) - c_n\big)\big(1 - \lambda\bar{F}_\omega(\sigma_{n,i})\big) \\ & + v_{i+2,\omega}\big(1 - \lambda\bar{F}_\omega(\sigma_{n,i})\big)\big(1 - \lambda\bar{F}_\omega(s^*)\big). \end{aligned} \tag{3.B.8}$$

Consider the strategy σ' in which the seller picks s^* in period i, picks s in period $i+1$, and otherwise acts as under σ. The contribution of periods before i as well as the probability of reaching period i under σ' is clearly the same as under σ. Conditional on period i being reached in state ω, the probability that period $i+2$ is reached under σ' is $\big(1 - \lambda\bar{F}_\omega(s^*)\big)\big(1 - \lambda\bar{F}_\omega(\sigma_{n,i})\big)$; this is the same as under σ. The continuation $v_{i+2,\omega}$ from period $i+2$ onwards in state ω is also unchanged by the deviation since her behaviour in periods $i+2$ onwards does not change. Thus we may evaluate the profit from the deviation in state ω by comparing the expression in (3.B.8) to the following:

$$\begin{aligned} & u_n(s^*)\lambda\bar{F}_\omega(s^*) \\ & + \big(u_n(\sigma_{n,i})\lambda\bar{F}_\omega(\sigma_{n,i}) - c_n\big)\big(1 - \lambda\bar{F}_\omega(s^*)\big) \\ & + v_{i+2,\omega}\big(1 - \lambda\bar{F}_\omega(s^*)\big)\big(1 - \lambda\bar{F}_\omega(\sigma_{n,i})\big). \end{aligned} \tag{3.B.9}$$

The deviation to σ' is profitable if (3.B.9) is strictly larger than (3.B.8). By rearranging, we find that deviation is profitable in state ω if and only if

$$u_n(s^*) - u_n(\sigma_{n,i}) \;>\; \frac{c_n}{\lambda}\frac{\underline{F}_\omega(s^*) - \underline{F}_\omega(\sigma_{n,i})}{\bar{F}_\omega(\sigma_{n,i})\bar{F}_\omega(s^*)}$$

holds. This inequality is implied by our choice of n^* and the assumption that $\sigma_{n,i}$ is a cutoff different from s^*. Thus the deviation to σ' is profitable in both states of the world, and so we have a contradiction to the fact that σ is a best response to $\pi_n^\emptyset$. □

Claim 3.18. *There exists an integer n^{**} such that if $n \geq n^{**}$, then $(\mu_n, \pi_n^\emptyset)$ is a sequential equilibrium of $V^\emptyset(n, c_n)$.*

Proof Claim 3.18. Recalling the construction of $(\mu_n, \pi_n^\emptyset)$ in Claim 3.15, it suffices to verify that μ_n is a best response to $\pi_n^\emptyset$ for all but finitely many n. Towards a contradiction, suppose not. Then there is a subsequence such that for each of its members there exists a profitable deviation from μ_n. By possibly relabelling, let this subsequence be the sequence itself. The expected utility of the seller is a linear function of her mixed strategy. Thus the assumption implies that for all n there exists a pure strategy σ_n such that

$$\begin{aligned} V^\emptyset(\sigma_n, \pi_n^\emptyset, n, c_n) \;&= \max_{\sigma'' \in S^n} V^\emptyset(\sigma'', \pi_n^\emptyset, n, c_n) \\ &> V^\emptyset(\mu_n, \pi_n^\emptyset, n, c_n). \end{aligned} \tag{3.B.10}$$

(The maximum is attained since S^n is finite.)

Let n^* be as in Claim 3.17. For integers n above n^*, we infer that there must exist i_n in $\{0, \ldots, n\}$ such that for all $i \in \{1, \ldots, n\}$ the following equivalence holds:

$$\sigma_{n,i} \;=\; s^* \qquad \Leftrightarrow \qquad i \;\leq\; i_n. \tag{3.B.11}$$

That is, the strategy σ plays s^* exactly up to some last period i_n, possibly never. In particular, we conclude that σ_n belongs to the set Σ_{n,i_n,s^*}.

Recall our construction of $(\mu_n, \pi_n^\emptyset)$. In particular, according to (3.B.7), we have

$$V^\emptyset(\mu_n, \pi_n^\emptyset, n, c_n) \;=\; \sup_{\sigma'' \in \Sigma_{n,j_n,s^*}} V^\emptyset(\sigma'', \pi_n^\emptyset, n, c_n). \tag{3.B.12}$$

Note that Σ_{n,j_n,s^*} contains Σ_{n,i,s^*} whenever i is an integer greater than j_n; for Σ_{n,i,s^*} contains exactly those pure strategies which play s^* for at least i periods, whereas the set Σ_{n,j_n,s^*} contains those strategies which play s^* for at least j_n periods. We have already argued that σ_n is in Σ_{n,i_n,s^*}. We therefore conclude from (3.B.10) that, for all n, the integer i_n from (3.B.11) is less than j_n.

Consider the pure strategy $\sigma_n^* = (s^*, \ldots, s^*)$, i.e. the strategy that constantly plays s^*. Note that σ_n^* is in Σ_{n,j_n,s^*}, so that (3.B.12) implies

$$V^\emptyset(\mu_n, \pi_n^\emptyset, n, c_n) \;\geq\; V^\emptyset(\sigma_n^*, \pi_n^\emptyset, n, c_n). \tag{3.B.13}$$

Using the inequality $i_n \le j_n$, we shall argue that for sufficiently large values of n we have

$$V^{\emptyset}(\sigma_n^*, \pi_n^{\emptyset}, n, c_n) \;>\; V^{\emptyset}(\sigma_n, \pi_n^{\emptyset}, n, c_n);$$

in light of (3.B.10) and (3.B.13), this yields a contradiction.

Consider the expected utility from σ_n^*, first. It is given by

$$\begin{aligned} & u_n(s^*) \sum_{\omega\in\{\ell,h\}} \sum_{i=1}^{n} \alpha_\omega \lambda \bar{F}_\omega(s^*) \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i-1} \\ & \qquad -\lambda c_n \sum_{\omega\in\{\ell,h\}} \sum_{i=1}^{n} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i-1} \\ = {} & u_n(s^*) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^{n}\right) \\ & \qquad -c_n \sum_{\omega\in\{\ell,h\}} \alpha_\omega \frac{1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^{n}}{\bar{F}_\omega(s^*)}. \end{aligned} \tag{3.B.14}$$

Now consider the expected utility from σ_n. We recall that σ_n selects s^* exactly up to some period i_n, where $i_n \le j_n$. By ignoring solicitation costs, we obtain an upper bound on the expected utility from σ_n. Verbally, a further upper bound is obtained in the following hypothetical scenario: If the seller does not trade within i_n periods, she gets a price of $\max_{s\in S\setminus\{s^*\}}(u_n(s))$ as soon she reaches period i_n+1. Formally,

$$\begin{aligned} & V^{\emptyset}(\sigma_n, \pi_n^{\emptyset}, n, c_n) \\ = {} & u_n(s^*) \sum_{\omega\in\{\ell,h\}} \sum_{i=1}^{i_n} \alpha_\omega \lambda \bar{F}_\omega(s^*) \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i-1} \\ & + \max_{s\in S\setminus\{s^*\}}(u_n(s)) \sum_{\omega\in\{\ell,h\}} \Bigg(\alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n} \\ & \qquad\qquad \times \sum_{i=i_n+1}^{n} \lambda\bar{F}_\omega(\sigma_{n,i}) \prod_{j=i_n+1}^{i-1} \left(1-\lambda\bar{F}_\omega(\sigma_{n,j})\right) \Bigg) \\ \le {} & u_n(s^*) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n}\right) \\ & + \max_{s\in S\setminus\{s^*\}}(u_n(s)) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n} \\ = {} & u_n(s^*) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n}\right) \\ & + \max_{s\in S\setminus\{s^*\}}(u_n(s)) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n} \\ = {} & u_n(s^*) + \left(\max_{s\in S\setminus\{s^*\}}(u_n(s)) - u_n(s^*)\right) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n}. \end{aligned} \tag{3.B.15}$$

We complete the argument by showing that (3.B.14) is greater than (3.B.15) for sufficiently large n. The difference (3.B.14) minus (3.B.15) is given by

$$
\begin{aligned}
&u_n(s^*) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^n\right) - c_n \sum_{\omega\in\{\ell,h\}} \alpha_\omega \frac{1-\left(1-\lambda\bar{F}_\omega(s^*)\right)^n}{\bar{F}_\omega(s^*)} \\
&- u_n(s^*) - \left(\max_{s\in S\setminus\{s^*\}} (u_n(s)) - u_n(s^*)\right) \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n} \\
&= \sum_{\omega\in\{\ell,h\}} \alpha_\omega \left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n} \Bigg(u_n(s^*) - \max_{s\in S\setminus\{s^*\}} (u_n(s)) - \frac{1}{\bar{F}_\omega(s^*)} \frac{c_n}{\left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n}} \\
&\qquad\qquad - \left(1-\lambda\bar{F}_\omega(s^*)\right)^{n-i_n} \left(u_n(s^*) - \frac{c_n}{\bar{F}_\omega(s^*)}\right)\Bigg)
\end{aligned}
$$

We know that $n-i_n$ is greater than $n-j_n$, and we know that the latter diverges as $n\to\infty$. Hence, for all ω, the term

$$\left(1-\lambda\bar{F}_\omega(s^*)\right)^{n-i_n} \left(u_n(s^*) - \frac{c_n}{\bar{F}_\omega(s^*)}\right)$$

converges to 0 as $n\to\infty$. We also recall that $u_n(s^*) - \max_{s\in S\setminus\{s^*\}}$ is positive and bounded away from 0 as $n\to\infty$. To prove that (3.B.14) minus (3.B.15) is strictly positive for sufficiently large n, it therefore suffices to show that

$$\frac{c_n}{\left(1-\lambda\bar{F}_\omega(s^*)\right)^{i_n}}$$

converges to 0 as $n\to\infty$. Again using that i_n is less than j_n, it suffices to check that

$$\frac{c_n}{\left(1-\lambda\bar{F}_\omega(s^*)\right)^{j_n}} \tag{3.B.16}$$

converges to zero. Recall that definition of j_n as

$$j_n = \min\left(\frac{1}{2}\frac{\ln v_\ell - \ln c_n}{-\ln(1-\lambda\bar{F}_h(s^*))}, \frac{n}{2}\right)$$

where we understand the minimum to be $n/2$ if $c_n = 0$. Therefore,

$$\frac{c_n}{\left(1-\lambda\bar{F}_\omega(s^*)\right)^{j_n}} \le c_n^{1/2} v_\ell^{1/2}.$$

Since $c_n \to 0$ as $n\to\infty$, we conclude that (3.B.16) converges to 0, as promised. □

In view of Claim 3.18, the next claim completes the proof.

Claim 3.19. *All of the following are true:*

(1) For all $n \in \mathbb{N}$, we have $\mu_n\{\sigma \in S^n : (\sigma_1, \ldots, \sigma_{j_n}) = (s^*, \ldots, s^*)\} = 1$.

(2) The sequence $\{j_n\}_{n \in \mathbb{N}}$ diverges to ∞.

(3) Along the sequence (μ_n, π_n), the good is traded with probability converging to 1. *The seller's expected utility and the price at which the good is traded converge almost surely to*

$$v_\ell + (v_h - v_\ell)\frac{\pi \frac{f_h(s^*)}{f_\ell(s^*)}\frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)}}{\pi \frac{f_h(s^*)}{f_\ell(s^*)}\frac{\bar{F}_\ell(s^*)}{\bar{F}_h(s^*)} + 1}. \tag{3.B.17}$$

Proof of Claim 3.19. Part (1) is immediate from the fact that μ_n is in Σ_{n,j_n,s^*}. Part (2) follows from the definition of j_n. Turning to part (3), note that the good is traded at a price of $u_n(s^*)$ whenever at least one of the first j_n buyers who arrives to the market has a signal equal to s^*. Conditional on state ω, the probability of this event is $\big(1 - \lambda f_\omega(s^*)\big)^{j_n}$. Since $u_n(s^*)$ converges to (3.B.17), we conclude from here that good is traded with probability converging to 1, and that the realized price conditional on trade converges almost surely to (3.B.17). It is clear that the seller's expected solicitation costs converge to 0, and hence the seller's expected utiltiy also converges to (3.B.17). □

□

3.B.4 Surplus extraction with binary signals

Proof of Proposition 3.7. Let $\bar{\sigma}_n$ denote the pure strategy that plays $\bar{s}$ in all periods. Given a strategy σ_n in S^n, we denote its i'th entry by $\sigma_{n,i}$.

Given $n \in \mathbb{N}$ and $m \in \{0, \ldots, n\}$, let $\sigma_n^{(m)}$ be the strategy which plays $\bar{s}$ in all rounds up to and including round m, and which plays $\underline{s}$ in all later rounds. Let us also abbreviate $\bar{x}_\omega = 1 - \lambda f_\omega(\bar{s})$.

Claim 3.20. *For all $m \in \{0, \ldots, n\}$ we have $\hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m)}) \geq \hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)$.*

Proof Claim 3.20. We will show that $\hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m)}) \geq \hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m+1)})$ holds for arbitrary $m \in \{1, \ldots, n-1\}$. This proves the claim since the strategy $\sigma^{(n)}$ is just the strategy $\bar{\sigma}_n$.

The difference $\hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m)}) - \hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m+1)})$ is given by

$$
\begin{aligned}
&\pi_0 \frac{\sum_{i=1}^{n} \mathbf{1}_{(\sigma_{n,i}^{(m)}=\bar{s})} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_h(\sigma_{n,j}^{(m)})\right)}{\sum_{i=1}^{n} \mathbf{1}_{(\sigma_{n,i}^{(m)}=\bar{s})} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_\ell(\sigma_{n,j}^{(m)})\right)} - \pi_0 \frac{\sum_{i=1}^{n} \mathbf{1}_{(\sigma_{n,i}^{(m+1)}=\bar{s})} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_h(\sigma_{n,j}^{(m+1)})\right)}{\sum_{i=1}^{n} \mathbf{1}_{(\sigma_{n,i}^{(m+1)}=\bar{s})} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_\ell(\sigma_{n,j}^{(m+1)})\right)} \\
=&\pi_0 \frac{\sum_{i=1}^{m} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_h(\sigma_{n,j}^{(m)})\right)}{\sum_{i=1}^{m} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_\ell(\sigma_{n,j}^{(m)})\right)} - \pi_0 \frac{\sum_{i=1}^{m+1} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_h(\sigma_{n,j}^{(m+1)})\right)}{\sum_{i=1}^{m+1} \prod_{j=1}^{i-1} \left(1-\lambda \bar{F}_\ell(\sigma_{n,j}^{(m+1)})\right)} \\
=&\pi_0 \frac{\sum_{i=1}^{m} (1-\lambda f_h(\bar{s}))^i}{\sum_{i=1}^{m} (1-\lambda f_h(\bar{s}))^i} - \pi_0 \frac{\sum_{i=1}^{m+1} (1-\lambda f_h(\bar{s}))^i}{\sum_{i=1}^{m+1} (1-\lambda f_h(\bar{s}))^i} \\
=&\pi_0 \frac{\sum_{i=1}^{m} \bar{x}_h^i}{\sum_{i=1}^{m} \bar{x}_\ell^i} - \pi_0 \frac{\sum_{i=1}^{m+1} \bar{x}_h^i}{\sum_{i=1}^{m+1} \bar{x}_\ell^i}.
\end{aligned}
$$

The sign of this difference is thus the sign of

$$
\begin{aligned}
&\left(\sum_{i=1}^{m} \bar{x}_h^i\right)\left(\bar{x}_\ell^{m+1} + \sum_{i=1}^{m} \bar{x}_\ell^i\right) - \left(\sum_{i=1}^{m} \bar{x}_\ell^i\right)\left(\bar{x}_h^{m+1} + \sum_{i=1}^{m} \bar{x}_h^i\right) \\
&= \sum_{i=1}^{m} \left(\bar{x}_h^i \bar{x}_\ell^{m+1} - \bar{x}_\ell^i \bar{x}_h^{m+1}\right).
\end{aligned}
$$

The claim now follows from the fact that $\bar{x}_\ell = 1 - \lambda f_\ell(\bar{s}) > 1 - \lambda f_h(\bar{s}) = \bar{x}_h$ holds. □

Claim 3.21. *Let $m \in \{1, \dots, n\}$. Let $\sigma \in S^n$. If σ_n is a permutation of $\sigma_n^{(m)}$, then $\hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n) \geq \hat{\pi}_n^{\emptyset}(\bar{s}, \sigma_n^{(m)})$.*

Proof Claim 3.21. For $k \in \{1, \dots m\}$, let $\iota_k(\sigma_n)$ denote the label of the round in which σ_n plays $\bar{s}$ for the k'th time.[19] Defining $\iota_k(\sigma_n^{(m)})$ analogously, notice that we have $\iota_k(\sigma_n^{(m)}) = k$.

Given a state ω, consider the sum

$$
\sum_{i=1}^{n} \mathbf{1}_{(\sigma_{n,i}=\bar{s})} \prod_{j=1}^{i-1} \left(1-\lambda(1-\underline{F}_\omega(\sigma_{n,i}))\right). \tag{3.B.18}
$$

The i'th summand is non-zero only if there is some k such that $i = \iota_k(\sigma_n)$ holds. In that case, the definition of $\iota_k(\sigma_n)$ implies the following: Over the course of rounds

19. That is, $\iota_1(\sigma_n) = \min\{i \in \{1, \dots, n\} : \sigma_{n,i} = \bar{s}\}$. The remaining indices are defined inductively via $\iota_k(\sigma_n) = \min\{i \in \{\iota_{k-1}(\sigma_n) + 1, \dots, n\} : \sigma_{n,i} = \bar{s}\}$.

$\{1,\ldots,\iota_k(\sigma_n)-1\}$, the strategy σ_n plays $\bar{s}$ exactly $k-1$ times, and $\underline{s}$ otherwise. Hence we have

$$\begin{aligned}
&\prod_{j=1}^{i-1}\big(1-\lambda(1-\underline{F}_\omega(\sigma_{n,i}))\big)\\
=&(1-\lambda(1-\underline{F}_\omega(\underline{s})))^{\iota_k(\sigma_n)-1-(k-1)}(1-\lambda(1-\underline{F}_\omega(\bar{s})))^{k-1}\\
&(1-\lambda)^{\iota_k(\sigma_n)-1-(k-1)}(1-\lambda f_\omega(\bar{s}))^{k-1}\\
=&(1-\lambda)^{\iota_k(\sigma_n)-1-(k-1)}\bar{x}_\omega^{k-1}.
\end{aligned}$$

The sum in (3.B.18) thus equals

$$\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_\omega^{k-1}.$$

A similar expression can be derived for $\sigma_n^{(m)}$, with the only change being that we have $\iota_k(\sigma_n^{(m)})=k$ for all k. The difference $\hat{\pi}_n^{\emptyset}(\bar{s},\sigma_n)-\hat{\pi}_n^{\emptyset}(\bar{s},\sigma_n^{(m)})$ thus reads

$$\begin{aligned}
&\pi_0\frac{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_h^{k-1}}{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_\ell^{k-1}}-\pi_0\frac{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n^{(m)})-k}\bar{x}_h^{k-1}}{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n^{(m)})-k}\bar{x}_\ell^{k-1}}\\
=&\pi_0\frac{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_h^{k-1}}{\sum_{k=1}^{m}(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_\ell^{k-1}}-\pi_0\frac{\sum_{k=1}^{m}\bar{x}_h^{k-1}}{\sum_{k=1}^{m}\bar{x}_\ell^{k-1}}.
\end{aligned}$$

The sign of this difference is the sign of

$$\begin{aligned}
&\sum_{k=1}^{m}\sum_{k'=1}^{m}\Big((1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_h^{k-1}\bar{x}_\ell^{k'-1}-(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_\ell^{k-1}\bar{x}_h^{k'-1}\Big)\\
=&\sum_{k,k'\,:\,k>k'}\Big((1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_h^{k-1}\bar{x}_\ell^{k'-1}-(1-\lambda)^{\iota_k(\sigma_n)-k}\bar{x}_\ell^{k-1}\bar{x}_h^{k'-1}\\
&\qquad+(1-\lambda)^{\iota_{k'}(\sigma_n)-k'}\bar{x}_h^{k'-1}\bar{x}_\ell^{k-1}-(1-\lambda)^{\iota_{k'}(\sigma_n)-k'}\bar{x}_\ell^{k'-1}\bar{x}_h^{k-1}\Big)\\
=&\sum_{k,k'\,:\,k>k'}\Big((1-\lambda)^{\iota_k(\sigma_n)-k}-(1-\lambda)^{\iota_{k'}(\sigma_n)-k'}\Big)\Big(\bar{x}_h^{k-1}\bar{x}_\ell^{k'-1}-\bar{x}_\ell^{k-1}\bar{x}_h^{k'-1}\Big)
\end{aligned}$$

To complete the proof, we argue that each of the summands in the last expression is weakly positive.

First, recall that $\iota_k(\sigma_n)$ denotes the label of the round in which σ_n plays $\bar{s}$ for k'th time, whereas $\iota_{k'}(\sigma_n)$ denotes label of the round with the (k')'th occurence. This means that at least $k-k'$ rounds must pass between the two rounds. Formally, we have $\iota_k(\sigma_n)-\iota_{k'}(\sigma_n)\geq k-k'$. This inequality implies that $(1-\lambda)^{\iota_k(\sigma_n)-k}-(1-\lambda)^{\iota_{k'}(\sigma_n)-k'}$ is negative.

Second, notice that $1-\lambda f_\ell(\bar{s}) = \bar{x}_\ell > \bar{x}_h = 1-\lambda f_h(\bar{s})$ holds. Given that the summands consider k and k' such that $k > k'$ holds, we conclude that $\bar{x}_h^{k-1}\bar{x}_\ell^{k'-1} - \bar{x}_\ell^{k-1}\bar{x}_h^{k'-1}$ is negative.

The previous two paragraphs imply that

$$\left((1-\lambda)^{\iota_k(\sigma_n)-k} - (1-\lambda)^{\iota_{k'}(\sigma_n)-k'}\right)\left(\bar{x}_h^{k-1}\bar{x}_\ell^{k'-1} - \bar{x}_\ell^{k-1}\bar{x}_h^{k'-1}\right)$$

is weakly positive, which yields the desired conclusion. □

Claim 3.22. *If $\mu_n' \in \Delta(S^n)$ is a mixed strategy that plays $\bar{s}$ with non-zero probability, then $\hat{\pi}_n^\emptyset(\bar{s}, \mu_n')$ is well-defined and we have $\hat{\pi}_n^\emptyset(\bar{s}, \mu_n') \geq \hat{\pi}_n^\emptyset(\bar{s}, \bar{\sigma}_n)$.*

Proof of Claim 3.22. For a pure strategy σ_n, a state ω and an integer i, let us abbreviate $\delta_\omega(i, \sigma_n) = \prod_{j=1}^{i-1}\big(1-\lambda\bar{F}_\omega(\sigma_{n,j})\big)$. In this notation, the sign of the difference $\hat{\pi}_n^\emptyset(\bar{s}, \mu_n') - \hat{\pi}_n^\emptyset(\bar{s}, \bar{\sigma}_n)$ is

$$\begin{aligned}
&\operatorname{sgn}\left(\frac{\sum_{\sigma\in S^n}\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\mu_n'(\sigma_n)\delta_h(i,\sigma_n)}{\sum_{\sigma\in S^n}\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\mu_n'(\sigma_n)\delta_\ell(i,\sigma_n)} - \frac{\sum_{i=1}^n \delta_h(i,\bar{\sigma}_n)}{\sum_{i=1}^n \delta_\ell(i,\bar{\sigma}_n)}\right)\\
&= \operatorname{sgn}\Bigg(\sum_{\sigma\in S^n}\mu_n'(\sigma_n)\Bigg(\left(\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\delta_h(i,\sigma_n)\right)\left(\sum_{i=1}^n \delta_\ell(i,\bar{\sigma}_n)\right)\\
&\qquad - \left(\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\delta_\ell(i,\sigma_n)\right)\left(\sum_{i=1}^n \delta_h(i,\bar{\sigma}_n)\right)\Bigg)\Bigg).
\end{aligned}$$

Hence it suffices to show that, for arbitrary σ_n, the difference

$$\left(\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\delta_h(i,\sigma_n)\right)\left(\sum_{i=1}^n \delta_\ell(i,\bar{\sigma}_n)\right) - \left(\sum_{i=1}^n \mathbf{1}_{(\sigma_{n,i}=\bar{s})}\delta_\ell(i,\sigma_n)\right)\left(\sum_{i=1}^n \delta_h(i,\bar{\sigma}_n)\right)$$

is weakly positive. There is nothing to prove if σ_n never plays $\bar{s}$. If σ_n plays $\bar{s}$ at least once, then the sign of this difference is precisely the sign of

$$\hat{\pi}_n^\emptyset(\bar{s}, \sigma_n) - \hat{\pi}_n^\emptyset(\bar{s}, \bar{\sigma}_n).$$

A strategy σ_n which plays $\bar{s}$ a total of, say, m times is a permutation of the strategy $\sigma_n^{(m)}$. Thus Claim 3.20 and Claim 3.21 together imply $\hat{\pi}_n^\emptyset(\bar{s}, \sigma_n) \geq \hat{\pi}_n^\emptyset(\bar{s}, \sigma_n^{(m)}) \geq \hat{\pi}_n^\emptyset(\bar{s}, \bar{\sigma}_n)$. In particular, $\hat{\pi}_n^\emptyset(\bar{s}, \sigma_n) - \hat{\pi}_n^\emptyset(\bar{s}, \bar{\sigma}_n)$ is weakly positive, as promised. □

Claim 3.23. *The seller's equilibrium expected utility converges to $\hat{v}_0$*

Proof of Claim 3.23. Recall that, for all n, the pair $(\mu_n, \pi_n^{\emptyset})$ is a sequential equilibrium. Claim 3.22 therefore implies that $\pi_n^{\emptyset}(\bar{s}) \geq \hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)$ holds for all n. It follows that the deviation to $\bar{\sigma}_n$ yields an expected utility of at least

$$\begin{aligned}
&\sum_{i=1}^{n} \sum_{\omega \in \{\ell,h\}} \alpha_\omega \left(\prod_{j=1}^{i-1} \left(1 - \lambda \bar{F}_\omega(\sigma_{n,j})\right) \right) \left(\lambda f_\omega(\bar{s}) \hat{v}(\bar{s}, \hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)) - c_n\right) \\
=&\sum_{i=1}^{n} \sum_{\omega \in \{\ell,h\}} \alpha_\omega \left(1 - \lambda f_\omega(\bar{s})\right)^{i-1} \left(\lambda f_\omega(\bar{s}) \hat{v}(\bar{s}, \hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)) - c_n\right) \\
=&\hat{v}(\bar{s}, \hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)) \sum_{\omega \in \{\ell,h\}} \alpha_\omega \left(1 - (1 - \lambda f_\omega(\bar{s}))^n\right) - c_n \sum_{\omega \in \{\ell,h\}} \alpha_\omega \frac{1 - \left(1 - \lambda f_\omega(\bar{s})\right)^n}{f_\omega(\bar{s})}
\end{aligned}$$

We know from Lemma 3.14 that $\hat{\pi}_n^{\emptyset}(\bar{s}, \bar{\sigma}_n)$ converges to $f_\ell(\bar{s})/f_h(\bar{s})$. Hence the expression in the previous line converges to $\hat{v}(\bar{s}, \pi_0 f_\ell(\bar{s})/f_h(\bar{s}))$ as $n \to \infty$. This expectation equals $\hat{v}_0$. Thus we have shown that equilibrium utility admits a lower bound which converges to $\hat{v}_0$. But we also know from Lemma 3.12 that equilibrium expected utility is bounded above by $\hat{v}_0$, and hence we arrive at the desired conclusion. □

□

References

Aliprantis, Charalambos D., and Kim C. Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide.* Third Edition. Springer. [94, 96]

Barsanetti, Bruno, and Braz Camargo. 2022. "Signaling in dynamic markets with adverse selection." *Journal of Economic Theory,* 105558. [93]

Bogachev, Vladimir I. 2007. *Measure theory.* Vol. 2. Springer Science & Business Media. [96]

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. 2006. "Dynamic monopoly pricing and herding." *RAND Journal of Economics* 37 (4): 910–28. [93, 94]

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. 2008. "Monopoly pricing in the binary herding model." *Economic Theory* 37 (2): 203–41. [93, 94]

Bose, Subir, Gerhard Orosel, and Lise Vesterlund. 2002. "Optimal pricing and endogenous herding." https://doi.org/https://dx.doi.org/10.2139/ssrn.318342. [93]

Chaves, Isaias N. 2019. "Privacy in bargaining: The case of endogenous entry." *Available at SSRN 3420766,* https://doi.org/https://dx.doi.org/10.2139/ssrn.3420766. [94]

Daley, Brendan, and Brett Green. 2012. "Waiting for News in the Market for Lemons." *Econometrica* 80 (4): 1433–504. [92]

Dilmé, Francesc. 2022. "Repeated bargaining with imperfect information about previous transactions." https://www.francescdilme.com/. [94]

Fuchs, William, Aniko Öry, and Andrzej Skrzypacz. 2016. "Transparency and distressed sales under asymmetric information." *Theoretical Economics* 11 (3): 1103–44. [77, 92]

Hörner, Johannes, and Nicolas Vieille. 2009. "Public vs. private offers in the market for lemons." *Econometrica* 77 (1): 29–69. [77, 92]

Hwang, Ilwoo. 2018. "Dynamic trading with developing adverse selection." *Journal of Economic Theory* 176: 761–802. [92]

Hwang, Ilwoo, and Fei Li. 2017. "Transparency of outside options in bargaining." *Journal of Economic Theory* 167: 116–47. [93]

Kaya, Ayca, and Qingmin Liu. 2015. "Transparency and price formation." *Theoretical Economics* 10 (2): 341–83. [92]

Kaya, Ayça, and Kyungmin Kim. 2018. "Trading dynamics with private buyer signals in the market for lemons." *Review of Economic Studies* 85 (4): 2318–52. [92]

Kaya, Ayça, and Santanu Roy. 2022a. "Market screening with limited records." *Games and Economic Behavior* 132: 106–32. [92]

Kaya, Ayça, and Santanu Roy. 2022b. "Price Transparency and Market Screening." https://doi.org/https://dx.doi.org/10.2139/ssrn.3662404. [92]

Kaya, Ayça, and Santanu Roy. 2022c. "Repeated Trading: Transparency and Market Structure." [92]

Kim, Kyungmin. 2017. "Information about sellers' past behavior in the market for lemons." *Journal of Economic Theory* 169: 365–99. [77, 92]

Krasteva, Silvana, and Huseyin Yildirim. 2012. "On the role of confidentiality and deadlines in bilateral negotiations." *Games and Economic Behavior* 75 (2): 714–30. [93]

Lauermann, Stephan, and Asher Wolinsky. 2016. "Search with adverse selection." *Econometrica* 84 (1): 243–315. [92, 93]

Pei, Harry. 2022a. "Building Reputations via Summary Statistics." https://doi.org/https://doi.org/10.48550/arXiv.2207.02744. [94]

Pei, Harry. 2022b. "Reputation Building under Observational Learning." *Review of Economic Studies.* [94]

Taylor, Curtis R. 1999. “Time-on-the-market as a sign of quality.” *Review of Economic Studies* 66 (3): 555–78. [77, 93]

Zhu, Haoxiang. 2012. “Finding a good price in opaque over-the-counter markets.” *Review of Financial Studies* 25 (4): 1255–85. [92]