Sequential Mechanisms for Evidence Acquisition¹

Elchanan Ben-Porath 2

Eddie Dekel³

Barton L. Lipman⁴

First Draft June 2023

¹We thank the National Science Foundation, grant SES–1919319 (Dekel and Lipman), and the US–Israel Binational Science Foundation for support for this research. We thank Simon Board, Duarte Goncalves, and Axel Niemeyer for useful comments and suggestions.

²Department of Economics and Center for Rationality, Hebrew University. Email: benporat@math.huji.ac.il.

³Economics Department, Northwestern University, and School of Economics, Tel Aviv University. Email: dekel@northwestern.edu.

⁴Department of Economics, Boston University. Email: blipman@bu.edu.

Abstract

We consider optimal mechanisms for inducing agents to acquire costly evidence in a setting where a principal has a good to allocate that all agents want. We show that optimal mechanisms are necessarily sequential in nature and have a threshold structure. Agents with higher costs of obtaining evidence and/or worse distributions of value for the principal are asked for evidence later, if at all. We derive these results in part by exploiting the relationship between the Lagrangian for this problem and the classic Weitzman (1979) "Pandora's box" problem.

1 Introduction

A principal has a single unit of a good or other resource to allocate to one of N agents under uncertainty regarding the value he would receive from allocating it to any given one of them. Each agent wants the good, independently of the value her receiving it would provide the principal. Each agent can obtain information which would reveal to her the value she would provide the principal if she receives the good as well as evidence which would prove this value to the principal. However, this information is costly to the agent, so she would not be willing to get it without being promised a sufficiently high chance that this will lead to her receiving the good. We characterize optimal mechanisms for the principal in this setting.

As examples of this situation, consider the head of an organization with multiple divisions, such as a university with multiple departments. The head of the organization has discrete resources to allocate, such as prestigious projects or assignments, or, in the case of a university, job slots. Some divisions would use this resource in ways that are more productive for the organization than others, but all divisions prefer to receive it. For a division to determine what value it would produce for the organization if it receives the resource is costly in time and/or effort, so that the division may prefer not to make a serious proposal to receive the resource.

We assume monetary transfers cannot be used. Intuitively, divisions of an organization have funds available to them to carry out actions on behalf of the organization. So it would be counterproductive to the organization to have divisions "bid" for these resources.

What, then, can the principal do? First, assume agents are symmetric both in their costs of obtaining evidence and in the probability distribution over the value they would provide the principal with the resource. Let $c \in (0, 1)$ denote the cost, let v_i denote a typical realization of the value to the principal of giving the good to agent *i* where these variables are iid across agents, and normalize the value of receiving the good to an agent to 1.

A natural mechanism to consider is for the principal to ask all agents to provide evidence, awarding the good to the agent who proves the highest value. If the number of agents, N, is large, this will not induce all agents to seek evidence. Specifically, if c > 1/N, the cost exceeds the expected benefit. Alternatively, the principal could exclude some agents from the mechanism and only ask some smaller number, say n, such that $1/n \ge c$. We will see that this can never be optimal.

In this symmetric problem, the optimal mechanism is easy to describe. The principal chooses a random ordering of the agents where all orderings are equally likely. He goes to the first agent in the selected ordering and asks her to provide evidence. In equilibrium, she will pay the cost c, learn her value, and prove this value to the principal. If her value is above a certain threshold, v^* , he stops and gives her the good. Otherwise, he continues to the next agent, applying the same threshold to decide whether to give her the good or continue. If all agents have values below v^* , he will end up asking all of them for evidence. In this case, he gives the good to the agent with the highest value.

An important point is that when the principal asks an agent for evidence, he does not tell her where she is in the sequence. To see why this is valuable, consider for simplicity the case of two agents. Suppose that when the agent is asked for evidence, she knows that she is second in line. Then she knows that the other agent's value is below the threshold. Letting F denote the common cdf for v, her probability of getting the good is then $1 - F(v^*) + (1/2)F(v^*) > 1/2$. This is because she certainly gets it if her value is above v^* and she is symmetric to the other agent conditional on her value being below v^* and so gets the good with probability 1/2 in this event. On the other hand, if she knows she is the first one to be asked, then her probability of getting the good must be strictly smaller. More specifically, it is $1 - F(v^*) + (1/2)[F(v^*)]^2$. This is because she knows that if she's below the threshold, the second agent will be asked for evidence and she'll only have a chance of getting the good if the other agent is also below the threshold. In short, the second agent has a larger incentive to pay for evidence than the first. Since the agents are symmetric, it is optimal to equalize the incentives by randomizing 50-50over which agent is asked first, rather than to distort the allocation differently for the two.

When the agents are not symmetric, several new considerations arise. First, because the principal has to give the good with a certain probability to any agent he asks for evidence, agents with higher probabilities on high values are better to ask earlier. Second, agents with higher costs have to be given more incentive to induce them to obtain evidence and hence he may wish to ask them later.

The optimal mechanism changes in three ways. First, instead of comparing the agent's values to the threshold or to one another, we compare virtual values. Specifically, for each agent *i*, instead of comparing v_i to the threshold or another agent's value, we compare $v_i + \lambda_i$ where λ_i reflects the severity of *i*'s incentive compatibility constraint. In fact, λ_i is the Lagrange multiplier on this constraint: agents who are harder to incentivize are given a constant "advantage" in the form of "extra points" added to their values.

Second, not surprisingly, the randomization over orders is no longer uniform. Indeed, in some cases, the asymmetries across agents will lead to some or all aspects of the ordering being deterministic. Intuitively, if one agent has much lower costs than another, then that agent can more easily bear the burden of being asked for evidence first.

Finally, if the asymmetries are significant enough, the mechanism exhibits a more

extreme form of ordering which we refer to as the use of *tiers*. More specifically, the agents are partitioned into tiers. The principal starts with the highest tier, asking agents in some random order for evidence. In each case, he compares the agent's virtual value to the tier 1 threshold, stopping and giving the good to the agent if her virtual value is above the threshold, continuing otherwise. If all the tier 1 agents have virtual values below the tier 1 threshold, the principal will learn all of their values. At this point, he then *lowers* the threshold. If any of the agents have virtual values above this lower threshold, he gives the good to the one with the highest virtual value. If not, he continues to the tier 2 agents now using this lower threshold for the tier 2 threshold. In the most extreme case, each agent is in her own tier and the order in which agents are asked for evidence is deterministic.

Intuitively, the use of tiers is valuable to the principal if the agents are extremely asymmetric. Suppose, for example, that the agents all have the same distribution of values and all but one have the same costs, with the remaining agent having a much higher cost than the others. Suppose the principal uses a mechanism with only one tier, with threshold v^* . Then even if the agent with the very high costs is last, we may need to make the threshold v^* very low and/or make her λ_i very high to give her an incentive to get evidence. Either is very costly in terms of the other agents. If we make v^* very low, then we are more likely to stop before getting to this agent and give the good to a relatively low type of another agent. If we make λ_i very large, then if none of the agents have virtual values above the threshold, the high–cost agent has such a large advantage that she is likely to get the good even when there is an agent with a significantly higher value. With the use of tiers, the principal can keep the threshold for the first N-1 agents at v^* and only bring it down if none of these agents has a high enough value. This way, the principal gets the value of the lower threshold in incentivizing the high–cost agent but loses less on the first N-1 agents than he would with a single threshold.

In principle, an optimal mechanism could also differ from the description above in one more way. Specifically, it could be that the randomization over the next agent to ask for evidence depends nontrivially on the result of evidence received from previous agents. We call such a mechanism a *generalized tiered threshold mechanism*, or a generalized mechanism for short. We use the term *tiered threshold mechanism* or, more briefly, a simple mechanism for the class of mechanisms described above where the random order is not conditional in this way. We show that if a generalized mechanism is optimal, then there is a simple mechanism which is incentive compatible and yields the principal and every type of every agent the same expected payoff. In this sense, there is no loss in restricting attention to simple mechanisms.

A key step in proving that the optimal mechanism takes this form is the exploitation of a formal connection to Weitzman's (1979) Pandora box problem. As we explain in more detail below, the Lagrangian for our problem takes exactly the form of Weitzman's problem if we treat the Lagrange multipliers as exogenous "preference parameters." This allows us to easily use Weitzman's characterization to show that every optimal mechanism is a generalized tiered threshold mechanism. Characterizing the multipliers and showing that we can restrict attention to simple mechanisms then completes the characterization of optimal mechanisms.

In Section 2, we lay out the details of the model. In Section 3, we characterize the optimal mechanism. In Section 4, we analyze the properties of the optimal mechanism, characterizing, in particular, the optimal random ordering of the agents. This section also provides some comparative statics results. We discuss some extensions of the model in Section 5. Finally, we offer brief concluding remarks in Section 6.

In the remainder of this section, we discuss the related literature. In addition to Weitzman (1979), there are three distinct branches of the related literature. First, our work is connected to the literature on mechanism design with evidence. Starting with Green and Laffont (1986), a number of papers consider properties of optimal mechanisms in settings where agents have access to hard information or evidence. See, for example, Glazer and Rubinstein (2004, 2006), Bull and Watson (2007), Deneckere and Severinov (2008), Hart, Kremer, and Perry (2017), and Ben Porath, Dekel, and Lipman (2019). In these papers, evidence is entirely exogenous: the agent simply has certain evidence as a function of her type. By contrast, here, no agent has evidence unless she pays a cost to obtain it.

A second literature considers mechanism design or games where agents can choose whether to acquire evidence. For example, Che and Kartik (2009) and DeMarzo, Kremer, and Skrzypacz (2019) consider game-theoretic models where an agent may choose a test which can generate an informative random signal. If obtained, this signal realization can be shown to another party, thus proving the realization and providing evidence. These papers do not consider mechanism design. By contrast, Ball and Kattwinkel (2022) and Ben Porath, Dekel, and Lipman (2022) do consider mechanism design when evidence can be acquired. These papers give some broad characterizations related to optimal mechanisms in these settings, but do not characterize optimal mechanisms for specific settings, as we do here.

Finally, there is a literature on mechanism design with information acquisition — see, for example, the survey of Bergemann and Välimäki (2006). In our model, the agent does not know her value until she acquires evidence, so evidence acquisition and information acquisition go hand-in-hand. The key difference between our work and these models, then, is exactly that the nature of the incentives to reveal the information acquired differ. With evidence, an agent is restricted in the misreports she can potentially get away with, so the honest-reporting constraints are different than in a model without evidence.

Perhaps the most closely related papers to ours are two papers in this literature,

namely, Gershkov and Szentes (2009) and Crémer, Spiegel, and Zhang (2009). Both can be described as settings with a principal and multiple agents. Gershkov–Szentes differs from our model in that they consider a *public* decision rather than a private allocation. The principal chooses a decision in $\{0, 1\}$ where all agents have the same state–dependent preferences between these options. In the optimal mechanism, the principal approaches agents in a random order to ask them to obtain information and provide it to him. Agents bear a private cost of obtaining information, as in our model. Because signals constitute evidence in our model but not in theirs, they need to impose truth–telling constraints in addition to the obedience constraints in both models. They consider only the case where agents are symmetric and restrict attention to mechanisms that are expost efficient. We do not need either of these restrictions.

Crémer, Spiegel, and Zhang, like us, consider a model where the principal is, in effect, allocating a single unit of a good. Unlike us, however, they do allow monetary transfers, which are critical to the model. Also unlike us, the principal does not inherently care which agent he gives the good to — instead, the principal is interpreted as a seller who maximizes the revenue he receives. In their model, the information acquisition by an agent is how the agent learns her valuation for the good. Crémer, Spiegel, and Zhang assume that the principal controls the information acquisition and so can block an agent from getting information before the principal is ready for her to do so. Essentially, they construct a VCG mechanism which extracts the entire ex ante surplus from the agents. Since agents can't get information before the principal is willing to let them, this is feasible. Then the principal pays the agents their information cost when he is ready for them to get information. In effect, the mechanism turns into a search problem where the principal seeks the most efficient way to find a high–value buyer to sell the good to. Because of this, they can also use Weitzman's (1979) results to characterize the optimal mechanism.

2 Model

There are $N \ge 2$ agents and a principal. The principal has one unit of a good to allocate to an agent. The value to the principal of allocating the good to agent *i* is v_i . However, v_i is unknown to the principal or any agent at the outset. We assume that the common prior over v_i is given by cdf F_i with strictly positive density f_i over the support [0, 1]. We assume v_i 's are independently distributed across agents. Aside from the assumption that the supports are the same, we impose no symmetry conditions on the distributions across agents. We sometimes let $V_i = [0, 1]$ denote the support of v_i and $V = [0, 1]^N = \prod_i V_i$.

Agent i can learn her value v_i at a cost $c_i \in (0, 1)$.¹ If she pays this cost, she not only

¹It is not difficult to extend our results to allow some agents to have costs above 1. See Section 5.

learns the realization of v_i but also obtains evidence enabling her to prove this realization to the principal. One can interpret "learning v_i " as observing a verifiable signal which generates a certain conditional expectation of v_i . Agent *i* can then prove this conditional expectation by showing this signal to the principal. Since all agents are risk neutral, replacing v_i with this conditional expectation changes nothing.

We assume all agents wish to obtain the good. Not including the cost of evidence acquisition, the agent's payoff is 1 if she receives the good, 0 otherwise. Her final payoff then is the payoff from the allocation of the good minus c_i if she acquired evidence. The principal's payoff is independent of whether/which agents incur evidence costs and is equal to 0 if he keeps the good and v_i if he gives the good to agent *i*.

In some examples, it is natural to assume that agent i cannot "consume" the good without paying cost c_i . For example, if we think of departments in a University competing for a job slot, then each department presumably must incur costs to identify their preferred candidate even if they are simply given the slot without needing information to compete for it. We assume that paying the cost is *not* necessary for consumption in this sense, but only because this is the more difficult case. As we explain in Section 5, one of our lemmas holds trivially if we assume it is necessary to pay this cost and no other results are affected.

The set of dynamic mechanisms available to the principal is quite complex. At each step, the principal can decide which agent or agents to ask for evidence, what (if anything) to tell them about what has happened so far, and how to react to the evidence they provide, if they do so. To keep the notation relatively simple, we restrict the class of mechanisms in a few ways that are clearly without loss of optimality for the principal.

First, we assume that the principal never asks more than one agent for evidence at a time. Because there is no discounting in our model, this is without loss for the principal as he can always ask one agent, then immediately afterward ask another.

Second, we assume that if an agent is asked to get evidence and refuses, then she is not given the good. By the Revelation Principle, we know that it is without loss of generality to consider mechanisms which induce agents to obey. Hence we may as well focus attention on mechanisms where agents are punished as severely as possible if they refuse to obey. In our model, the most severe possible punishment for an agent is not giving her the good.

Third, we assume no agent is asked for evidence more than once. As noted above, the Revelation Principle tells us we can focus on mechanisms in which agents always obey. Since the principal cannot gain by getting evidence from an agent twice, he will never ask any agent more than once. Finally, we assume that the principal never provides any information for an agent upon asking her to obtain evidence. Put differently, each agent has (at most) one information set in the game the principal induces. To see that this is without loss for the principal, suppose instead that there are two different information sets for the agent. Then incentive compatibility requires that the agent's expected utility to obeying the principal is higher than her expected payoff to disobeying conditional on each information set. This implies but is not implied by the statement that the agent's expected utility to obeying is higher than her expected utility to disobeying unconditionally. Hence the principal weakly improves incentives by pooling these histories together.

Intuitively, a dynamic mechanism specifies one of the following decisions by the principal after every history. Either (a) the principal ends the process and keeps the good, (b) the principal ends the process and gives the good to one of the agents, or (c) the principal asks some agent to obtain and provide evidence.

More formally, a history is a sequence of agents and their responses to being asked for evidence. To be specific, a length n history is a sequence $((i_1, x_1), (i_2, x_2), \ldots, (i_n, x_n))$ with the following properties. First, $i_n \in \{1, \ldots, N\}$ for all n. That is, i_n is the agent who is the nth agent asked for evidence. Second, $x_n \in [0, 1] \cup \{R\}$. Here $x_n = R$ denotes the response of agent i_n to refuse to provide evidence. If $x_n \in [0, 1]$, then x_n is the value proved by agent i_n . Finally, $i_k = i_\ell$, then $k = \ell$ — that is, no agent can be asked for evidence twice.

Let H_n denote the set of all length n histories and let $H = \bigcup_{n=0}^N H_n$ where we define $H_0 = \{e\}$, so that e is the empty history. Note that there cannot be a history of length more than N.

A (pure) dynamic mechanism is a measurable function $d: H \to (\{0\} \times \{0, 1, ..., N\}) \cup (\{1\} \times \{1, ..., N\})$ satisfying the properties stated below. If d(h) = (0, i), this means the principal ends the process on history h and gives the good to agent i (where i = 0 — that is, keeping the good — is possible). If d(h) = (1, i), this means the principal continues the process on history h and asks agent i for evidence. Note that d(h) cannot equal (1, 0) — that is, the principal cannot ask himself for evidence.

The properties we require d to satisfy are as follows. First, we require

$$d((i_1, x_1), (i_2, x_2), \dots, (i_n, x_n)) \neq (0, i_k)$$

if $x_k = R$ for any k = 1, ..., n. That is, if i_k was asked for evidence and refused, she cannot get the good.

Second, we require that

$$d((i_1, x_1), (i_2, x_2), \dots, (i_n, x_n)) \neq (1, i_k)$$

for any k = 1, ..., n. That is, the principal cannot ask any agent for evidence who has already been asked.

Let \hat{D}_p denote the set of pure dynamic mechanisms and let \hat{D} denote the set of probability mixtures over \hat{D}_p — that is, the set of random mechanisms.

The following lemma shows that we can restrict the set of dynamic mechanisms further. We will restrict attention to mechanisms that satisfy a property we call *no free lunch*. We say a mechanism satisfies no free lunch if it never gives the good to an agent who has not provided evidence. It is easy to see that the principal may as well choose a mechanism satisfying this property as he does not pay the costs. The following lemma establishes the more substantial point that *every* optimal dynamic mechanism satisfies this property.

Lemma 1. Fix any incentive compatible dynamic mechanism violating no free lunch on a set of histories with strictly positive probability. Then there is another incentive compatible mechanism which gives the principal a strictly higher expected payoff. Hence any optimal incentive compatible mechanism must satisfy no free lunch up to sets of measure zero.

Proof. Fix a set of histories H^* with positive probability on which the principal gives the good to agent i with positive probability even though i has not provided evidence. There are two cases to consider. First, suppose that every history in H^* has the property that the principal has previously received evidence from at least one agent. (If H^* does not have this property but some positive measure subset does, we can replace H^* with this subset.) For each history $h \in H^*$, let $\bar{v}(h)$ denote the highest value which some other agent has previously proven to the principal. Because this set of histories has positive measure, the set of values of $\bar{v}(h)$ for this set of histories must be uncountable. Hence the probability that v_i and h have the property that $v_i < \bar{v}(h)$ must be strictly positive.

Let \bar{v}_i be defined by $1 - F_i(\bar{v}_i) = c_i$. Because $c_i < 1$ for all *i* and because F_i is continuous for all *i*, we know that $\bar{v}_i > 0$. Fix some $\varepsilon \in (0, \bar{v}_i)$.

Change the mechanism only on histories in H^* as follows. With whatever probability the original mechanism gave the good to agent i, the principal instead asks agent i to provide evidence. If i does not provide evidence, the principal keeps the good. If i does provide evidence and shows either that $v_i \geq \bar{v}_i - \varepsilon$ or that $v_i \geq \bar{v}(h)$, the principal gives the good to i. If the agent's evidence shows that $v_i < \bar{v}_i - \varepsilon$ and $v_i < \bar{v}(h)$, the principal gives the good to the agent who has proven value $\bar{v}(h)$.

It is easy to see that i's probability of receiving the good if she provides evidence is strictly larger than her cost, so it is strictly optimal for her to provide evidence.² Also,

²Note that it is optimal for this agent to get evidence whether the principal informs her that she is

every other agent's incentive to provide evidence is at least as large as before since any agent who obtains evidence now gets rewarded with the good more often. Hence the new mechanism is incentive compatible. Clearly, the probability that the principal gains by giving the good to an agent with a value higher than v_i is strictly positive, so the principal's expected payoff to the alternative mechanism is strictly larger. Hence the original mechanism was not optimal.

Second, suppose that the set of histories for which the principal gives the good to i with positive probability without obtaining evidence from i has positive probability, but the subset of these histories on which some other agent has previously provided evidence has probability zero. By incentive compatibility, this means that on these histories, no other agent has been asked for evidence. In this case, fix any agent $j \neq i$ and replace $\bar{v}(h)$ in the argument above with $E(v_j)$. The change in the mechanism now has no effect on the incentives of other agents to obtain evidence since the change only takes place when none of them have been asked to do so. Hence the new mechanism is again incentive compatible and improves the principal's payoff, implying that the original mechanism was not optimal.

Let D_p denote the set of pure dynamic mechanisms $d \in \hat{D}_p$ satisfying no free lunch i.e., those such that d(h) = (1, i) only if $h = ((i_1, x_1), (i_2, x_2), \dots, (i_n, x_n))$ where $i = i_k$ for some k. Let D denote the set of probability mixtures over D_P .

By the Revelation Principle, it is without loss of generality to focus attention on mechanisms in which agents find it optimal to obey the principal. That is, once we restrict to incentive compatible mechanisms, we know that the relevant histories will be ones where agents who are asked for evidence do, in fact, provide it. Given such a dynamic mechanism, we can then compute the outcome under the mechanism as a function of the profile of types v.³

Let $P(d) = (P_1(d), \ldots, P_N(d))$ denote the allocation probabilities induced by dynamic mechanism d. That is, for each d, $P_i(d)$ is a measurable function mapping V to [0, 1] where

$$\sum_{i} P_i(v \mid d) \le 1, \quad \forall v \in V.$$

Let $e_i(d)$ denote the probability agent *i* is asked for evidence in mechanism *d*.

on one of these histories or if the principal pools this set of histories with any other histories on which she would be asked for evidence in the original mechanism. In the latter case, we are pooling two sets of histories, where it is optimal for the agent to get evidence conditional on either separately and hence conditional on the union.

³We omit the precise definition as it will not be needed. Briefly, one can iteratively define the probability distribution over realized histories given d as a function of the profile v. This then determines the probability distribution over outcomes.

We can now state the principal's maximization problem. The principal's objective function is

$$\mathbf{E}_{v}\left[\sum_{i} P_{i}(v \mid d)v_{i}\right].$$

The constraints are the N incentive compatibility constraints. One might expect these constraints to be very complex since they say that *conditional on being asked for evidence*, an agent finds it optimal to obey. This conditioning depends in a complex way on the dynamic mechanism since we have to identify the set of histories on which this agent might be asked for evidence. Fortunately, we are able to bypass this complexity by expressing the incentive compatibility constraint at the ex ante stage. Recall that each agent has (at most) one information set in the mechanism. Hence we can write the incentive compatibility constraint as requiring that the ex ante optimal strategy for the agent is to obtain evidence and provide it if asked for it. Note that if the agent obeys the principal, then her expected payoff in mechanism d must be

$$\mathbf{E}_{v}P_{i}(v \mid d) - c_{i}e_{i}(d).$$

To compare this to the agent's expected payoff to any ex ante strategy which disobeys, note first that if the agent is not asked for evidence, then, by the no-free-lunch restriction, she will not get the good. Hence regardless of her deviation strategy, her payoff if not asked for evidence is 0.

To pin down the agent's deviation payoff more precisely, first, consider the deviation strategy where the agent does get evidence when asked but does not always report it. In this case, her expected costs of evidence acquisition are still $c_i e_i(d)$, but she must receive the good weakly less often. Hence obeying the principal must give a weakly higher payoff than this. Second, consider the deviation strategy of not getting evidence. In this case, she gets a payoff of 0 whether she is asked for evidence or not and hence her expected payoff is 0. Therefore, we can write the incentive compatibility constraint for agent *i* as

$$\mathcal{E}_v P_i(v \mid d) \ge c_i e_i(d).$$

Hence we can state the principal's optimization problem as follows. We say that $d \in D^*$ is *optimal* if it solves the problem

$$\max_{d \in D} \mathcal{E}_{v} \left[\sum_{i} P_{i}(v \mid d) v_{i} \right]$$

subject to

$$\mathbf{E}_{v}[P_{i}(v \mid d)] - c_{i}e_{i}(d) \ge 0, \quad \forall i.$$

Let D^* denote the set of optimal d's. In the next section, we characterize this set.

3 Characterizing the Optimal Mechanism

We show that without loss of utility for the principal, we can focus on a class of mechanisms we call *tiered threshold mechanisms* or sometimes *simple mechanisms* for short. A tiered threshold mechanism consists of the following. We have a partition of the set of agents into K tiers, denoted $\mathcal{I}_1, \ldots, \mathcal{I}_K$. For each tier \mathcal{I}_k , we have two additional objects. First, we have a random ordering of the agents in that tier. More specifically, there is a probability distribution, denoted O_k over the set of linear orders over \mathcal{I}_k where we interpret a typical such order, \succ_k , by saying that if $i \prec_k j$, then i goes before j. Second, for each tier \mathcal{I}_k , we have a threshold $v_k^* \in \mathbf{R}_+$. Finally, for each agent i, we have a non-negative number λ_i . We refer to $v_i + \lambda_i$ as i's virtual value.

Given these objects, the tiered threshold mechanism works as follows. First, we draw random orders over the sets of agents in each tier. Let \succ_k be the ordering drawn for tier k. We then ask the first agent according to \succ_1 for evidence. (In equilibrium, all agents obey requests for evidence.) If $v_i + \lambda_i > v_1^*$ so that her virtual value is above the tier 1 threshold, the mechanism stops and gives her the good. Otherwise, we continue to the next agent according to \succ_1 . If all the tier 1 agents have virtual values below v_1^* , then all will be asked for evidence. At this point, if any of these agents has a virtual value above the tier 2 threshold, v_2^* , the mechanism stops and gives the good to that agent with the highest virtual value. If not, we continue to the first agent in tier 2 according to \succ_2 . Again, the mechanism stops and gives the good to this agent if her virtual value is above the tier 2 threshold and continues otherwise. If none of the tier 2 agents has a virtual value in this range, we compare the virtual values of all tier 1 and tier 2 agents to the tier 3 threshold v_3^* , etc.

If all agents have virtual values below the tier K threshold, then the mechanism will end up asking all of them for evidence. In this case, the good is allocated to that agent with the highest virtual value.

We also show that the class of *all* optimal mechanisms is the set of what we call *generalized tiered threshold mechanisms* or generalized mechanisms for brevity. The only difference between a generalized mechanism and a simple mechanism is in the distribution over orders within a tier. In a simple mechanism, there is a single random choice of an order associated with each tier. In a generalized mechanism, which agent within a tier is chosen at any ponit can depend on all past observations by the principal.

Our proofs connect the dynamic mechanism design problem to Weitzman's (1979) "Pandora's box" problem. First, we briefly summarize Weitzman's results.

Weitzman considers the following problem, simplified here to more easily line up with our problem. There is a searcher who faces N "boxes." Box *i* has a certain monetary prize x_i in it where x_i is distributed according to a distribution \hat{F}_i . Prizes are independently distributed across boxes. There is a cost, \hat{c}_i , to opening box *i*. At each point in the search process, the searcher decides between quitting and taking no box, quitting and taking some box she has previously opened, or opening a box she has not yet opened. The searcher's payoff is the prize in the box she takes (or 0 if she takes no box) minus the accumulated costs of the boxes she has opened.

Weitzman characterizes the set of optimal search procedures as follows. For each box i, associate an index, r_i , defined by

$$\hat{c}_i = \mathcal{E}_{x_i} \max\{x_i - r_i, 0\}.$$

Intuitively, r_i is that value such that if the searcher could quit and take r_i or open box i and then quit taking the better of r_i and the prize in box i, she would be indifferent between these options. For simplicity, we assume $r_i > 0$ for all i, as the analog of this property will necessarily hold for our use of Weitzman's result. Given this, a search procedure is optimal iff it takes the following form. First, the searcher opens any box i_1 with the highest index — i.e., such that $r_{i_1} = \max_j r_j$. If there is more than one such box, any randomization is optimal. If the prize in box i_1, x_{i_1} , satisfies $x_{i_1} > \max_{j \neq i_1} r_j$, then the searcher stops and takes box i_1 . In our problem, the analog of the x_i 's will be continuously distributed, so we do not need to consider what happens if $x_{i_1} = \max_{j \neq i_1} r_j$. If $x_{i_1} < \max_{j \neq i_1} r_j$, the searcher opens any box i_2 with the highest index of the remaining boxes — i.e., such that $r_{i_2} = \max_{j \neq i_1} r_j$. The searcher continues in this fashion, comparing the largest prize found so far in any box to the maximum index of the unopened boxes. If the largest prize is strictly above the highest index among the unopened boxes, the searcher stops and takes the corresponding box. Otherwise, she continues and opens any unopened box with the largest possible index. If she opens all boxes, she takes the one with the largest prize.

The following lemma is the key to linking Weitzman's result to our problem. Recall that D^* is the set of optimal d's. Given $\lambda \in \mathbf{R}^N_+$, let $D^{**}(\lambda)$ denote the set of maximizers of the Lagrangian function

$$\mathcal{L} = \mathcal{E}_{v} \left[\sum_{i} P_{i}(v \mid d) v_{i} \right] + \sum_{i} \lambda_{i} \left[\mathcal{E}_{v} P_{i}(v \mid d) - c_{i} e_{i}(d) \right]$$

and let D^{**} denote the set of incentive compatible $d \in D^{**}(\lambda)$ for some λ such that $\lambda_i[\mathbb{E}_v P_i(v \mid d) - c_i e_i(d)] = 0$ for all i.

Lemma 2. $D^* = D^{**}$.

In other words, strong duality holds. This follows from the fact that the set of (P, e) that can be generated by a mechanism is convex and that payoffs are linear in (P, e). The

proof of this result is relatively standard but is contained in the supplemental appendix for the convenience of the reader.

Given this lemma, Weitzman's result almost immediately yields the following:

Theorem 1. Every optimal dynamic mechanism is a generalized tiered threshold mechanism.

Proof. We can rewrite the Lagrangian as

$$\mathbf{E}_{v}\left[\sum_{i} P_{i}(v \mid d)(v_{i} + \lambda_{i}) - \lambda_{i}c_{i}e_{i}(d)\right].$$

Let d^* denote an optimal mechanism and λ^* the associated Lagrange multipliers. By Lemma 2, d^* solves the problem

$$\max_{d \in D} \mathbf{E}_{v} \left[\sum_{i} P_{i}(v \mid d)(v_{i} + \lambda_{i}^{*}) - \lambda_{i}^{*}c_{i}e_{i}(d) \right].$$

This is almost exactly Weitzman's problem. To be specific, think of agent i as box i where the prize in box i is $v_i + \lambda_i^*$ and the cost of opening box i is $\lambda_i^* c_i$. Think of $P_i(v \mid d)$ as the probability of choosing box i under search procedure d when the prizes are given by $(v_1 + \lambda_1^*, \ldots, v_N + \lambda_N^*)$ and $e_i(d)$ as the probability of opening box i under search procedure d.

The only difference between this problem and Weitzman's is that a mechanism in our problem must specify what to do on a history where some agent has been asked for evidence and refused to comply. In Weitzman's problem, a search procedure is not defined on such a history as boxes cannot refuse to be opened. Let H^* denote the set of all possible histories in our problem with the property that no agent has ever refused when asked for evidence. Then the set of search procedures in Weitzman is exactly our set of dynamic mechanisms when we restrict the set of histories to H^* . In both cases, the options are only to quit and receive a payoff of zero, quit and give the good to some agent who has provided evidence/take a previously opened box, or to ask some agent for evidence who has not yet provided it/open some unopened box. Because an agent will not refuse to provide evidence if asked, the histories we exclude when considering H^* are payoff irrelevant.

Hence d^* must be in the class of search procedures Weitzman identifies. The index for box/agent *i*, which we denote by $\hat{v}_i + \lambda_i^*$, is defined by

$$\lambda_i^* c_i = \mathcal{E}_{v_i + \lambda_i^*} \max\{v_i + \lambda_i^* - (\hat{v}_i + \lambda_i^*), 0\} = \mathcal{E}_{v_i} \max\{v_i - \hat{v}_i, 0\}.$$

Partition the set of agents into tiers, $\mathcal{I}_1, \ldots, \mathcal{I}_K$ where the agents in \mathcal{I}_1 have the largest index, those in \mathcal{I}_2 have the next largest index, etc. For tier \mathcal{I}_k , let v_k^* denote the common index for the agents in that tier.

Then the optimal procedure must start with the agents in tier 1, asking them for evidence in some order. This order is arbitrary in Weitzman and so could depend on the specific agents previously asked for evidence and/or the evidence they show. If the prize in box *i*/associated with agent *i* of $v_i + \lambda_i$ is larger than v_1^* , then we stop and take that box/give the good to that agent. Otherwise, we continue to some other box/agent in tier 1. After checking the last agent in tier 1, the relevant comparison is to the common index for the second tier, v_2^* . Thus if the agent with the highest virtual value in tier 1 is above v_2^* , this agent gets the good and otherwise we continue to tier 2.

It is not hard to see that this is exactly a generalized tiered threshold mechanism.

As the proof of Theorem 1 shows, we use Weitzman's theorem to say that the optimal mechanism has the form of a generalized tiered threshold mechanism. However, Weitzman's results do not identify the randomization over the order of checking, not even whether it varies with previous observations by the principal. In Weitzman, if two boxes have the same index, then any randomization over which to check first is equally good, including randomizations that depend on the prizes in previously opened boxes. For our model, though, these randomizations are significant as they matter for incentive compatibility. As discussed in the introduction, all else equal, an agent who is asked earlier for evidence is less likely to receive the good. Hence an agent who is more likely to be asked early has less incentive to obey when asked for evidence. In other words, Weitzman's theorem identifies the form of the optimal mechanism for some profile of λ_i 's without identifying anything about the randomizations involved. The next step is to identify the λ_i 's and characterize the randomizations.

A key simplification is that we can restrict attention to the simpler class of tiered threshold mechanisms, rather than considering all generalized tiered threshold mechanisms. Intuitively, the application of Weitzman's theorem tells us that the payoff of the principal is not directly affected by the randomizations, so the principal does not directly gain from making these depend on the past history. As we will show, having these randomizations depend on the past history does not help with incentive compatibility either, so there is no value to them.

More specifically, for any generalized tiered threshold mechanism, we will show that there is a simple mechanism which is equivalent in the following sense.

Definition 1. Mechanisms d and d' are interim–equivalent if for all i and all $v_i \in [0, 1]$, $E_{v_{-i}}P_i(v \mid d) = E_{v_{-i}}P_i(v \mid d')$ and if for all i, $e_i(d) = e_i(d')$. Because the agents are not endowed with private information, it would arguably be appropriate to call two mechanisms equivalent if they gave the principal and every agent the same ex ante expected payoff. We use the stronger notion of interim–equivalence both because it gives a stronger result and because the interim comparison is a convenient approach to proving our results.

For brevity, we define the usual interim probabilities (also called the reduced form) by $p_i(v_i \mid d) = \mathbb{E}_{v_{-i}} P_i(v_i, v_{-i} \mid d)$.

The following lemma shows the significance of interim–equivalence.

Lemma 3. If d is an optimal incentive-compatible mechanism and d' is interim-equivalent to it, then d' is also an optimal incentive-compatible mechanism and every type of every agent obtains the same payoff in d' as in d.

Proof. Since d is incentive compatible, we have

$$\mathbb{E}P_i(v \mid d) - c_i e_i(d) \ge 0$$

or

$$\mathbf{E}_{v_i} p_i(v_i \mid d) - c_i e_i(d) \ge 0$$

for all *i*. Hence, since d' is interim–equivalent to d, we have

$$\mathbf{E}_{v_i} p_i(v_i \mid d') - c_i e_i(d') \ge 0$$

for all i, so d' is also incentive compatible. More generally, the payoff to agent i of type v_i in mechanism d is

$$E_{v_{-i}}P_i(v_i, v_{-i} \mid d) - c_i e_i(d) = p_i(v_i \mid d') - c_i e_i(d'),$$

so every type of every agent is indifferent between the two mechanisms.

Finally, we can write the payoff of the principal under d as

$$\mathbf{E}_{v}\left[\sum_{i} P_{i}(v \mid d)v_{i}\right] = \sum_{i} \mathbf{E}_{v_{i}}\left[\mathbf{E}_{v_{-i}}P_{i}(v_{i}, v_{-i} \mid d)v_{i}\right] = \sum_{i} \mathbf{E}_{v_{i}}\left[p_{i}(v_{i} \mid d)v_{i}\right].$$

Since d' is interim–equivalent to d, the principal's payoff is the same under d and d', so if d is optimal, d' must be optimal as well.

Put differently, Lemma 3 says that we can only identify optimal mechanisms (at most) up to interim–equivalence. Hence we may as well focus on a convenient selection from the interim–equivalent optimal mechanisms. We show below that tiered threshold mechanisms are such a selection.

This claim will follow from a result which is broadly useful for characterizing the optimal mechanism. Specifically, we show that we can identify the λ_i 's and the e_i 's and use these variables to identify the randomizations. In particular, any two mechanisms which have the same λ_i 's and the same e_i 's are interim–equivalent. We then characterize the set of feasible e_i 's given the λ_i 's and show that any feasible e_i can be generated by a simple tiered threshold mechanism, implying that we can restrict attention to these mechanisms.

Recall that a tiered threshold mechanism specifies numbers $\lambda_1, \ldots, \lambda_N$, a partition of the agents into tiers $\mathcal{I}_1, \ldots, \mathcal{I}_K$, and thresholds for each tier. The difference between generalized and simple tiered threshold mechanisms is how the order of asking agents within a tier is determined. As the proof of Theorem 1 shows, once we specify the λ_i 's, the indices $\hat{v}_i + \lambda_i$ are defined by Weitzman's formula, which in turn defines the tiers and thresholds. Hence the λ_i 's and the order of asking agents within each tier is all that needs to be specified.

Lemma 4. Fix mechanisms d and \hat{d} in $D^{**}(\lambda)$ satisfying $e_i(d) = e_i(\hat{d})$ for all i. Then d and \hat{d} are interim–equivalent.

Proof. We prove this result by showing that we can write $p_i(v_i \mid d)$ entirely as a function of the λ 's and e_i . Given this, if two mechanisms have the same λ 's and e's, they must be interim–equivalent.

So fix any $\lambda \in \mathbf{R}^N_+$ and any $d \in D^{**}(\lambda)$. The proof of Theorem 1 shows that this must be a generalized tiered threshold mechanism. Fix the *e*'s generated by this mechanism.

The proof of Theorem 1 shows that we can define the index for *i* entirely from λ_i . Specifically, it is $\hat{v}_i + \lambda_i$ where \hat{v}_i is defined by

$$\lambda_i c_i = \int_{\hat{v}_i}^1 (v - \hat{v}_i) f_i(v) \, dv.$$

It is easy to show that \hat{v}_i is uniquely determined by λ_i .

Given the profile $(\hat{v}_1 + \lambda_1, \dots, \hat{v}_N + \lambda_N)$ and $e = (e_1, \dots, e_N)$, we can compute $p_i(v_i \mid d)$ for any *i* and any v_i as follows. First, if $v_i \geq \hat{v}_i$, then agent *i* receives the good if and only if she is asked for evidence. Second, if $v_i < \hat{v}_i$, we have

$$p_i(v_i \mid d) = \prod_{j \neq i \mid \hat{v}_j + \lambda_j \ge v_i + \lambda_i} F_j(v_i + \lambda_i - \lambda_j).$$

To see this, first, consider $v_i \geq \hat{v}_i$. Clearly, the no-free-lunch property implies that *i* does not receive the good if she is not asked for evidence. If she is asked for evidence and is in a tier with other agents, she receives the good if her virtual value is above the

threshold for that tier. Since the threshold is the common index for the agents in that tier, this means she receives the good if $v_i + \lambda_i \ge \hat{v}_i + \lambda_i$ or simply $v_i \ge \hat{v}_i$. If she is in a tier by herself, she receives the good if her virtual value is above the next highest index among the other agents. But then if her virtual value is above her own index, this certainly holds. Hence if $v_i \ge \hat{v}_i$, we have $p_i(v_i \mid d) = e_i(d)$.

So suppose $v_i < \hat{v}_i$, so *i*'s virtual value is below her index. In this case, *i* receives the good only if all the other agents who are checked have virtual values below hers. To be more specific, note that any agent *j* with $\hat{v}_j + \lambda_j > v_i + \lambda_i$ will be checked before agent *i* is given the good. So for *i* to receive the good when her value is v_i , it must be the case that all such *j* have $v_j + \lambda_j < v_i + \lambda_i$. The expression above gives exactly the probability of this event.

Hence $p_i(v_i \mid d)$ is uniquely identified by the λ 's and e's generated by d.

The next step is to identify the set of e_i 's that can be generated by an optimal mechanism given the λ 's. To see the issue, suppose that tier 1 consists of agents 1 and 2, that $\hat{v}_1 = \hat{v}_2 \in (0, 1)$, and that $e_1 = e_2 = 1$. From the above, if $v_1 \geq \hat{v}_1$ and $v_2 \geq \hat{v}_2$, both agents 1 and 2 have virtual values above the threshold for tier 1. Hence both get the good iff they are asked for evidence. But we are hypothesizing that both are asked for evidence with probability 1. This is simply not possible — when $v_1 \geq \hat{v}_1$ and $v_2 \geq \hat{v}_2$, we are saying both agents get the good, but only one of them can.

In other words, given the λ 's, not every (e_1, \ldots, e_N) can be generated by some choice of randomization over the order of asking agents for evidence. The issue may seem different, but it turns out to be related to Border's (1991) characterization of the set of interim allocation functions which are feasible in the sense that they can be generated by some allocation functions. Border's result covered symmetric distributions and subsequent work extended his results in many directions — see, for example, Mierendorff (2011) or Che, Kim, and Mierendorff (2013). Here we give some clearly necessary conditions on the e_i 's in the spirit of the Border conditions. In the Appendix, we show that these conditions are sufficient.

Lemma 5. Fix $\lambda = (\lambda_1, \ldots, \lambda_N)$ and the associated $\hat{v}_1, \ldots, \hat{v}_N$. If there exists a generalized tiered threshold mechanism $d \in D^{**}(\lambda)$ which generates $(e_1(d), \ldots, e_N(d)) = (e_1, \ldots, e_N)$, then the following conditions hold. First,

$$\sum_{i \in \mathcal{I}_k} e_i [1 - F_i(\hat{v}_i)] = \left[\prod_{i \in \mathcal{I}^{k-1}} F_i(v_k^* - \lambda_i)\right] \left[1 - \prod_{i \in \mathcal{I}_k} F_i(\hat{v}_i)\right],\tag{1}$$

where the first term on the right-hand side is defined to be 1 for k = 1 and otherwise

$$\mathcal{I}^{k-1} = \bigcup_{\ell=1}^{k-1} \mathcal{I}_{\ell}.$$

Second, for all $\mathcal{I} \subset \mathcal{I}_k$,

$$\sum_{i \in \mathcal{I}} e_i [1 - F_i(\hat{v}_i)] \le \left[\prod_{i \in \mathcal{I}^{k-1}} F_i(v_k^* - \lambda_i) \right] \left[1 - \prod_{i \in \mathcal{I}} F_i(\hat{v}_i) \right],\tag{2}$$

Furthermore, for any e satisfying these conditions, there is a simple tiered threshold mechanism $d \in D^{**}(\lambda)$ that generates $e_i(d) = e_i$ for all i.

Because of their similarity to the conditions identified by Border (1991), we refer to equations (1) and (2) as the Border conditions.

Proof. The proof that any e generated by a generalized tiered threshold mechanism must satisfy the Border conditions is straightforward. First, consider tier 1. For k = 1, the first Border condition, equation (1), says

$$\sum_{i \in \mathcal{I}_1} e_i [1 - F_i(\hat{v}_i)] = 1 - \prod_{i \in \mathcal{I}_1} F_i(\hat{v}_i).$$

To understand the left-hand side, consider an agent i in tier 1 with $v_i \geq \hat{v}_i$ (equivalently, $v_i + \lambda_i \geq z_i^*(\hat{v}_i)$). By the definition of a generalized tiered threshold mechanism, such an agent gets the good if and only if she is asked for evidence. Hence the left-hand side is the probability that the good goes to some agent i in tier 1 with a value in this range. However, the definition of a generalized tiered threshold mechanism also says that if there is some agent i in tier 1 with a value in this range, then the good *must* go to such an agent. Since the right-hand side is the probability that the realized v has at least one tier 1 agent with a value in this range, we see that the left-hand side and right-hand side must be equal.

Continuing with k = 1, the second Border condition, equation (2), says

$$\sum_{i \in \mathcal{I}} e_i [1 - F_i(\hat{v}_i)] \le 1 - \prod_{i \in \mathcal{I}} F_i(\hat{v}_i),$$

for all $\mathcal{I} \subseteq \mathcal{I}_1$. To see that this must hold, note that the left-hand side is the probability that an agent $i \in \mathcal{I}$ has $v_i \geq \hat{v}_i$, is asked for evidence, and hence receives the good, while the right-hand side is the probability that some agent $i \in \mathcal{I}$ has $v_i \geq \hat{v}_i$. Hence the left-hand side must be smaller than the right.

Similarly, consider k = 2, where the first Border condition says that

$$\sum_{i \in \mathcal{I}_2} e_i[1 - F_i(\hat{v}_i)] = \left| \prod_{i \in \mathcal{I}_1} F_i(v_2^* - \lambda_i) \right| \left| 1 - \prod_{i \in \mathcal{I}_2} F_i(\hat{v}_i) \right|$$

Now the left-hand side is the probability that an agent *i* in tier 2 has a value $v_i \geq \hat{v}_i$, is asked for evidence, and hence receives the good. We know that this happens if and only if all the tier 1 agents have virtual values below the tier 2 threshold v_2^* and some tier 2 agent has $v_i \geq \hat{v}_i$. A tier 1 agent *i* has virtual value below the tier 2 threshold iff $v_i + \lambda_i < v_2^*$ or $v_i < v_2^* - \lambda_i$. So the right-hand side is exactly the probability a tier 2 agent with virtual value above the tier 2 threshold gets the good.

For k = 2, the second Border condition says that

$$\sum_{i \in \mathcal{I}} e_i [1 - F_i(\hat{v}_i)] \le \left[\prod_{i \in \mathcal{I}_1} F_i(v_2^* - \lambda_i) \right] \left[1 - \prod_{i \in \mathcal{I}} F_i(\hat{v}_i) \right],$$

for every $\mathcal{I} \subseteq \mathcal{I}_2$. In this case, the left-hand side is the probability that some agent $i \in \mathcal{I}$ has a value above \hat{v}_i and gets the good, while the right-hand side is the obviously necessary condition that no agent in tier 1 gets the good before we get to tier 2 and that there is some agent $i \in \mathcal{I}$ with $v_i \geq \hat{v}_i$. Hence this inequality is necessary.

The proof for other tiers is analogous. The proof that any e satisfying the Border conditions can be generated by a simple tiered threshold mechanism is in the Appendix.

Corollary 1. For any optimal generalized tiered threshold mechanism d, there is a simple tiered threshold mechanism d' which is interim–equivalent to it. Hence there is always an optimal mechanism which is a tiered threshold mechanism.

To see why the corollary holds, note that Lemma 5 implies that the e generated by any $d \in D^{**}(\lambda)$ can also be generated by a simple tiered threshold mechanism $\hat{d} \in D^{**}(\lambda)$. By Lemma 4, then, for any generalized tiered threshold mechanism, there is an interimequivalent simple tiered threshold mechanism.

Summarizing, we have shown

Theorem 2. An incentive-compatible mechanism d is optimal if and only if it is a generalized tiered threshold mechanism with tiers defined by $\hat{v}_1 + \lambda_1, \ldots, \hat{v}_N + \lambda_N$ where $\lambda_i \geq 0$ for all i,

$$\lambda_i c_i = \int_{\hat{v}_i}^1 (v - \hat{v}_i) f_i(v) \, dv, \ \forall i,$$
$$\lambda_i \left[\mathbf{E}_{v_i} p_i(v_i \mid \lambda, e_i) - c_i e_i \right] = 0, \ \forall i$$

and the Border conditions (1) and (2) hold, where $e_i = e_i(d)$ for all *i*.

Furthermore, given any optimal mechanism d, there is a tiered threshold mechanism which is also optimal and yields every type of every agent the same expected payoff. A tempting but incorrect intuition suggests that the complexities of identifying these randomizations can "typically" be avoided. We only need to identify the random order in which agents are asked when agents are in the same tier. Agents are only in the same tier when they have the same index. It is tempting to suspect that for "generic" c_i 's and F_i 's, indices never tie, so, in this sense, randomization is "almost always" irrelevant.

This is not correct for the following reason. Recall that the index is $\hat{v}_i + \lambda_i$ where \hat{v}_i is defined by

$$\lambda_i c_i = \int_{\hat{v}_i}^1 (v - \hat{v}_i) f_i(v) \, dv.$$

Evidently, the index depends not just on the exogenous c_i and F_i , but also on the endogenous λ_i . This endogeneity leads to ties in the indices and hence with more than one agent with nontrivial probability. In fact, "similar enough" agents *must* be in the same tier, so that ties are not "measure zero."

To see the intuition, first consider the (nongeneric) symmetric agent case. Suppose there are two agents with the same F_i 's and the same c_i 's. For simplicity, assume the incentive compatibility constraint is binding for both agents. In this case, the optimal mechanism *must* randomize 50–50 over who to start with. The key reason for this is that, as is easily shown, $\hat{v}_i + \lambda_i$ is decreasing in λ_i in the relevant range. So suppose we try to construct an optimal mechanism where these two agents are in different tiers. Without loss of generality, suppose we try to put 1 in the higher tier, so $\hat{v}_1 + \lambda_1 > \hat{v}_2 + \lambda_2$. The functions defining \hat{v} from λ are the same for the two agents since they have the same c_i and F_i . So the fact that the index is decreasing in λ implies that we must have $\lambda_1 < \lambda_2$.

But then 1 has the pressure of going first and the disadvantage of a smaller "handicap" in the form of a smaller λ . Given that the agents are identical, this implies that the incentive constraint cannot be binding for both agents.⁴

So suppose $\hat{v}_1 + \lambda_1 = \hat{v}_2 + \lambda_2$, so the two agents are in the same tier, but we randomize over which is first in some asymmetric fashion. Because the two agents have the same index, they must have the same λ_i 's. Thus everything that enters their payoffs is the same *except* the randomization over which is asked for evidence first. Since the one more likely to go first must have a lower expected payoff, the only way both incentive constraints

⁴To see this, note that if both incentive constraints are binding, then the total expected evidence costs must equal 1 as the total probability of allocating the good is 1. Because 1 is always asked for evidence since she is first, her expected evidence cost is c, so 2's must be 1-c. Hence 2 must be asked for evidence with probability (1-c)/c. Therefore, 1 gets the good without 2 being asked with probability 1-(1-c)/c. That is, $F(\hat{v}_1) = (1-c)/c$. Since 1's probability of getting the good must be c in total, this means that the probability 1 gets the good when 2 is asked is $c - [1 - (1-c)/c] = (1-c)^2/c$. But the handicap of the smaller λ means that 1 gets the good less than half of the time that both 1 and 2 are below the threshold. Hence this must be less than $(1/2)[(1-c)/c]^2$. But this holds iff $c \leq 1/2$, implying that the first-best is feasible and the incentive constraints do not bind.

can bind is if the randomization is 50–50.

Given that a 50–50 randomization is the unique solution for identical agents, it should not be surprising that nearby randomizations are the unique solution for nearly identical agents. Hence ties are not "nongeneric."

The figure below illustrates. Assuming two agents, each with $v_i \sim U[0, 1]$, the area between the red curves is the set of (c_1, c_2) in the range $[.5, 1]^2$ where the two agents have the same index in the optimal mechanism. As the intuition above suggests, it is a non-negligible set of types around symmetry.



4 Properties of the Optimal Mechanism

In this section, we describe properties of optimal mechanisms. Section 4.1 specializes the discussion to the particularly tractable symmetric case, describing the mechanism and its properties in detail. In Section 4.2, we return to the general case and characterize the random ordering of agents used in the optimal dynamic mechanism. Finally, Section 4.3 gives some comparative statics.

4.1 Symmetric Mechanisms

When the agents are symmetric in the sense that $c_i = c_j \equiv c$ and $F_i = F_j \equiv F$ for all *i* and *j*, the analysis simplifies greatly. In this case, the principal doesn't care *which* agent

he asks for evidence next and hence only needs to decide *whether* to seek evidence at any particular point.

More formally, it is not hard to see that in the symmetric setting, there must be an optimal mechanism which is symmetric in the sense that it treats the agents identically. Thus λ_i , \hat{v}_i , and e_i are all independent of i. For simplicity, we drop the i subscripts on these variables when discussing the symmetric case.

Since all agents have the same $\hat{v} + \lambda$, clearly, there is only one tier and one threshold. In this case, there is no point distinguishing between virtual values, $v_i + \lambda$, and actual values, v_i , since the λ will cancel out of any comparison across agents or comparison of an agent to the threshold. Hence we may as well simplify and ignore the λ .

In this case, the optimal mechanism is simply stated. If the incentive compatibility constraint does not bind, then $e_i = 1$ for every agent *i* and $\hat{v}_i = 1$. That is, every agent is asked for evidence and the good is allocated to the agent with the highest value. Hence incentive compatibility binds iff c > 1/N.

If the incentive compatibility constraint does bind, then $e_i = 1/(Nc)$ for all *i*. To see this, recall that the incentive compatibility constraint is that

$$\mathbf{E}_{v_i} p_i(v_i) \ge c_i e_i.$$

In the symmetric case, all agents are equally likely to receive the good, so the left-hand side is 1/N. Since this constraint binds and $c_i = c$ for all i, we see that $e_i = 1/(Nc)$ for all i. Finally, Lemma 5 implies that \hat{v} is pinned down by

$$\sum_{i} e_i [1 - F_i(\hat{v}_i)] = 1 - \prod_{i} F_i(\hat{v}_i)$$

or, using symmetry,

$$Ne[1 - F(\hat{v})] = 1 - [F(\hat{v})]^N$$

Using the formula for the sum of a geometric series, we see that this is equivalent to

$$e = \frac{1}{N} \sum_{j=0}^{N-1} [F(\hat{v})]^j.$$

To see how this lines up with the dynamic mechanism, recall that we choose an order at random with, in the symmetric case, all orders equally likely.⁵ This means that any given agent *i* has a 1/N chance of being first, 1/N chance of being second, etc. Think

⁵To be sure, there are other ways to generate the same evidence probabilities. For example, we could randomize uniformly over the orders (1, 2, ..., N). (2, 3, ..., N, 1), (3, 4, ..., N, 1, 2), ..., (N, 1, 2, ..., N-1).

of the index j on the right-hand side as denoting how many agents are ahead of i in the selected order. If j = 0, then i is first and hence asked for evidence with certainty. If j = 1, there is one agent ahead of i and so i is asked for evidence iff this agent has a value below \hat{v} . Hence in this case, i is asked for evidence with probability $F(\hat{v})$. In general, there are j agents ahead of i with probability 1/N and i is asked for evidence in this situation with probability $[F(\hat{v})]^j$, the probability all these agents have values below \hat{v} . In short, using the value computed earlier for e, we see that \hat{v} is defined by

$$\frac{1}{c} = \sum_{j=0}^{N-1} [F(\hat{v})]^j.$$

The comparative statics for the symmetric case are straightforward. If we increase c, the left-hand side of the equation above falls, so \hat{v} must be reduced. This is natural — if the cost of acquiring evidence goes up, agents must be promised a higher chance of getting the good to induce them to obtain evidence. This requires reducing the threshold. The evidence probability e also is reduced. Again, this fits together intuitively with the reduction in the threshold. With a lower threshold, each agent is less likely to be reached and asked for evidence.

If the distribution F of values is shifted up in the sense of first-order stochastic dominance, then $F(\hat{v})$ is smaller at every point. Hence we must increase \hat{v} to restore equality. Again, this is intuitive: if agents are more likely to have high values, then the principal can raise the threshold and be "pickier" about who he gives the good to. Note that e is unchanged since it is 1/(Nc). So the improvement in an agent's probability of getting the good from the improvement in F is entirely extracted by the principal in raising the threshold.

Finally, the effects of increasing the number of agents, N, is similarly straightforward to compute, though the intuition is less immediate. If we increase N, then we must reduce \hat{v} to restore equality in the equation above. So if there are more agents, the principal holds each to a *lower* standard, a perhaps unexpected conclusion. Intuitively, the principal does this because the increase in the number of agents reduces any one agent's likelihood of getting the good if the threshold is unchanged. Thus each agent's incentive to obtain evidence is reduced and must be restored by lowering the threshold. Increasing N also lowers the probability any given agent gets the good and lowers the probability the agent is asked for evidence.

One might expect that the increase in the number of agents and resulting lowering of the threshold could make the principal worse off. While it can be optimal for the principal to exclude some agents in asymmetric settings, this is never true in the symmetric case. Increasing the number of agents in the symmetric case always makes the principal strictly better off. One way to see this is to consider the expected number of agents asked for evidence. We know that if the principal asks, say, n agents for evidence, then each agent has an equal probability of being one of the n agents asked. Hence each agent has probability n/N of being asked for evidence in this case. So, overall, the expected probability that an agent is asked for evidence is 1/N times the expected number of agents asked. Since we know this probability is 1/(Nc), this says that the expected number of agents asked for evidence is 1/c, independent of N. On the other hand, if the number of agents goes from N to N+1, the probability the principal asks N+1 agents for evidence goes from 0 to something strictly positive. Because the threshold falls, the probability the principal asks only one agent for evidence also increases. In other words, as we increase the number of agents asked. This change is valuable to the principal. It enables him to sample more agents when the draws are all low, though at the cost of sampling fewer when the draws are high.

Another way to see the point is to suppose that the principal designs a mechanism for N symmetric agents and then one more is added. Suppose that the principal is restricted to asking at most N agents for evidence, but can give the good to the (N+1)st agent. Suppose the principal chooses N of the N+1 agents at random, with all agents equally likely to be included, and then runs the N agent mechanism with the chosen subset. If he ends up asking all N agents for evidence and all have values below E(v), then, instead of giving the good to the agent with the highest value, he gives it to the (N+1)st agent, the one he did not get evidence from. Clearly, if this is incentive compatible, it is better for the principal than the usual N agent mechanism. To see that it must be incentive compatible, consider any agent's probability of being asked for evidence. If the agent is one of the N agents chosen at the outset, then her probability is the same as in the N-agent mechanism. Of course, if she is not chosen, her probability of being asked for evidence is 0. Hence her overall probability of being asked is N/(N+1) times the probability of being asked in the usual N-agent mechanism. But we know that the probability of being asked in the N-agent mechanism is 1/(Nc), so in this mechanism, it is 1/[(N+1)c]. Hence her expected evidence cost is 1/(N+1). Because the overall mechanism treats all agents symmetrically, her probability of receiving the good is also 1/(N+1), so this mechanism is indeed incentive compatible. So the principal is strictly better off when N + 1 agents than with N.

It is worth noting that this observation generalizes to case where not all of the agents are symmetric in the following way.⁶

Theorem 3. Suppose $c_i = c_j$ and $F_i = F_j$. In any optimal mechanism, we have $e_i = 0$ iff $e_j = 0$.

⁶This result follows directly from Theorem 4 below, but it is more convenient to prove it separately.

4.2 Optimal Ordering

In the symmetric case, the agents are identical, so there is no reason for the principal to prefer one order for seeking evidence over another. When agents are asymmetric, what determines the optimal randomization over the order?

Intuitively, agents who are later in the order are more protected from competition. These later agents are only asked for evidence when the values of the earlier agents are relatively low, so being asked is a good sign for them about the competition they face. This suggests that the principal will tend to put "stronger" agents earlier. The theorem we present in this section shows that this intuition is exactly correct.

We say that agent *i* is stronger than agent *j* if $c_i \leq c_j$ and F_i (weakly) first-order stochastically dominates F_j . Note that one or both of these comparisons can be an equality relation — i.e., agents with the same cost and same distribution are each stronger than the other.

Theorem 4. If *i* is stronger than *j*, then in an optimal mechanism, we have $e_i \ge e_j$.

The following corollaries elucidate the implications of this result.

Corollary 2. If *i* is stronger than *j*, then *i*'s index, $\hat{v}_i + \lambda_i$ is weakly larger than *j*'s. Hence if *i* is stronger than *j* and they are in different tiers, then *i* is in a higher tier than *j*. In particular, if the optimal order is deterministic, *i* is asked before *j*.

To see why the corollary follows, first note that if i and j have different indices and hence are in different tiers, then $e_i \ge e_j$ implies that i must be the one in the higher tier. This is because agents in tier k are not asked for evidence until *all* agents in *all* higher tiers have been asked. Similarly, if the optimal order is deterministic, agents who are later in the order are necessarily asked for evidence with lower probability than those before.

Corollary 3. If *i* is stronger than *j* and *j* is stronger than *i*, then $e_i = e_j$. In this case, the optimal order cannot be deterministic if their incentive constraints are binding.

If *i* and *j* have the same costs and same distribution, each is stronger than the other, so Theorem 4 implies that $e_i = e_j$. If the order is deterministic, this means that *i* is asked for evidence if and only if *j* is also asked. This cannot be optimal if the incentive constraints for *i* and *j* are both binding, so in this case, the optimal order *cannot* be deterministic. Note that if both incentive constraints are slack, both agents are asked for evidence with probability 1 and the principal does not care whether he asks both at the same time or in some order.

Corollary 4. If *i* is stronger than *j* and *i* is excluded in the sense that she receives the good with zero probability, then *j* is excluded. This implies that in the symmetric case, exclusion is not optimal.

To see this, first observe that *i* is excluded if and only if $e_i = 0$. If $e_i = 0$, then the no-free-lunch property implies that *i* never receives the good and hence is excluded in this sense. Conversely, if *i* never receives the good, then her incentive constraint becomes $-c_i e_i \ge 0$, implying $e_i = 0$. If *i* is stronger than *j* and *i* is excluded, then we have $e_i = 0$ and hence $e_j = 0$.

In a symmetric model, all agents are stronger than all other agents. Hence all must have the same e_i . So if one of them is excluded in the optimal mechanism, all must be. But this can never be optimal since the principal could do better simply by giving the good to one of the agents without asking anyone for evidence.

4.3 Comparative Statics

As noted in Section 4.1, comparative statics in the symmetric case are straightforward to derive. Unfortunately, this is not true in the asymmetric case.

To see why the comparative statics are complex and can vary across the parameter space, consider the effects of an increase in c_1 in the two-agent case. Throughout this discussion, we assume both incentive constraints bind. First, suppose we start at a point where $\hat{v}_1 + \lambda_1 > \hat{v}_2 + \lambda_2$. In this case, the optimal mechanism begins by asking 1 for evidence. If $v_1 + \lambda_1 > \hat{v}_2 + \lambda_2$, or, equivalently, $v_1 > v^* \equiv \hat{v}_2 + \lambda_2 - \lambda_1$, 1 receives the good. Otherwise, 2 is asked for evidence and whichever agent has the higher virtual value receives the good. That is, 1 receives the good iff $v_1 + \lambda_1 > v_2 + \lambda_2$ or $v_1 > v_2 + \lambda_2 - \lambda_1$. Note, then, that the allocation of the good depends only on v^* and $\lambda_2 - \lambda_1$.

When c_1 increases, we must change the allocation or else 1's incentive constraint will be violated. Assume the change in c_1 is small so that we continue to have $\hat{v}_1 + \lambda_1 > \hat{v}_2 + \lambda_2$ and hence continue to start with agent 1. Then we must change v^* and $\lambda_2 - \lambda_1$ in such a way as to improve 1's probability of receiving the good to offset the increase in her costs without violating 2's incentive constraint.

It is not hard to see that the variables v^* and $\lambda_2 - \lambda_1$ affect 2 in opposite directions but affect 1 in the same direction. To be specific, decreases in v^* and in $\lambda_2 - \lambda_1$ both improve 1's payoff, while the first improves 2's payoff and the second reduces it.

To see this, consider first a reduction in v^* . Clearly, this reduction helps 1 since she is more likely to receive the good without having to compete with 2. To see that it also helps 2, simply note that the best types of agent 1 that 2 had to compete with are now receiving the good without 2 being asked for evidence. Hence, conditional on being asked, 2's chances of receiving the good are improved.

On the other hand, a reduction in $\lambda_2 - \lambda_1$ (holding v^* fixed) hurts 2 but helps 1. This is because 2's relative handicap in the competition with 1 is lowered, so, all else equal, 2 loses to 1 more often.

Because we need to maintain 2's incentive constraint, we must move v^* and $\lambda_2 - \lambda_1$ in the same direction at magnitudes such that the net effect on 2's payoff is zero. Clearly, then, we must reduce both: increasing both could not possibly compensate 1 for the increase in c_1 .

While this explanation focuses on the case where 1 is asked for evidence first, essentially the same argument applies to the situation where 2 is asked first.

Consider, by contrast, the situation where e_1 and e_2 are both in (0, 1). In this case, the optimal mechanism has a strictly interior probability of asking 1 for evidence first. Now the simple analysis above falls apart and a wide range of things can happen. This is simply because the principal now has another tool to compensate 1 for the increase in her costs. In addition to the kind of changes discussed above, the principal could now respond by lowering e_1 or, equivalently, lowering the probability that 1 is the first to be asked for evidence.

In this situation, the principal could, for example, lower the probability of starting with 1, which lowers e_1 and raises e_2 , lower the threshold v^* , and raise 2's relative advantage $\lambda_2 - \lambda_1$. The change in the e's helps 1 and hurts 2, the change in the threshold helps both, and the change in the relative advantage hurts 1 and helps 2. Hence by making these changes in the appropriate magnitudes, 2's utility could remain unchanged so that her incentive constraint continues to hold, while 1's utility is improved to offset the effect of the increase in c_1 .

In fact, it is not hard to show that when we start from symmetry, this is what must happen. That is, if $F_1 = F_2$ and we start from the point $c_1 = c_2$, we know that we start where each agent is first with probability 1/2. A small increase in c_1 from this point necessarily leads to a reduction in $\lambda_2 - \lambda_1$, exactly the opposite of what we see if we start from parameters where a deterministic order is optimal. Not only can $\lambda_2 - \lambda_1$ increase or decrease, one can also show that the threshold v^* can increase or decrease, and e_2 can increase or decrease.⁷

More broadly, with N agents and anywhere from 1 to N different tiers possible, the

⁷We conjecture that an increase in c_1 must weakly reduce e_1 , but even this is not clear.

principal has a range of tools which affect the agents in different directions. The value of these tools to the principal and for improving agent incentives vary in complex ways across the parameter space. As a result, there are essentially no comparative statics that hold globally.

5 Extensions

5.1 High Costs/Reserve Value

Up to this point, we have assumed $c_i < 1$ for all *i*. In this section, we describe how the analysis changes if some agents have $c_i \ge 1$.

It is easy to see that the principal cannot get (useful) information from agent i if $c_i \ge 1$. Obviously, if $c_i > 1$, then i strictly prefers not getting the good to getting it and providing evidence. If $c_i = 1$, the principal can induce i to obtain evidence, but only by promising to give i the good regardless of what her value turns out to be. In this sense, the principal can get information from i if $c_i = 1$ but cannot get information he can actually use.

So if there are such agents, their only role is as a kind of reservation value for the principal. In other words, we have assumed that the principal receives a payoff of 0 from keeping the good. If there is some agent i with $c_i \ge 1$, then the principal has a better outside option than keeping the good since he can give it to i for a "known" payoff of $E(v_i)$. Hence adding such agent simply changes the principal's reservation utility.

This has the same effect as adding an agent i with $c_i = 0$ and a known value. The optimal mechanism can be thought of as asking this agent for evidence first with $\lambda_i = 0$ and $\hat{v}_i = 1$. In other words, this agent is asked first but the principal continues after learning her type (except in the zero measure situation where her value is 1). Then this agent's value is compared to thresholds and other agents' types as above.

5.2 Varying Value to Agent from Receiving Good

The use of Weitzman's results to characterize the structure of the optimal mechanism enables us to extend the analysis to the case where the agent's value of receiving the good varies and may be correlated with the value to the principal of giving it to her.

For an illustration, suppose that if the value to the principal of giving the good to the

agent is v_i , then the value to agent *i* of receiving it is $\varphi_i(v_i)$. Assume $\varphi_i(\cdot)$ is continuous and that $E_{v_i}\varphi_i(v_i) = 1$. Then it is not hard to show that our result that every optimal mechanism satisfies no-free-lunch continues to hold. Now the Lagrangian takes the form

$$E_{v}\left[\sum_{i} P_{i}(v \mid d)v_{i} + \sum_{i} \lambda_{i} \left(P_{i}(v \mid d)\varphi_{i}(v_{i}) - e_{i}(d)c_{i}\right)\right]$$
$$E_{v}\left[\sum_{i} P_{i}(v \mid d)(v_{i} + \lambda_{i}\varphi_{i}(v_{i})) - \sum_{i} \lambda_{i}c_{i}e_{i}(d)\right].$$

or

As above, we can view the λ 's as fixed parameters and characterize the solution to this maximization problem using the Weitzman solution for the case where the prize in box i is $v_i + \lambda_i \varphi_i(v_i)$. In this problem, we see that if φ_i is decreasing, it is possible that agent i is more likely to receive the good for some low values of v_i than for some higher values.

5.3 Costs

If we change the principal's objective so that he prefers, all else equal, to not impose costs on the agents, then with no other changes in the model, we lose the no-free-lunch property and hence the ability to appeal to Weitzman (1979). On the other hand, we can still use Weitzman's result for a variation of our model.

To be specific, assume the principal's payoff if he allocates the good to agent i is $v_i - \sum_j \alpha_j c_j x_j$ where $x_j = 1$ if agent j gets evidence and is zero otherwise. Also, assume that agent i cannot consume the good unless she acquires evidence. For example, if we think of the agents as departments in a university competing for a job slot and the c_i 's as the costs of determining who they would hire, then it seems clear that agent i cannot use the job slot without acquiring information. In this case, it is without loss of optimality for the principal to restrict attention to mechanisms that satisfy no-free-lunch. If the principal is about to allocate the good to agent i without requiring evidence from her, he knows his payoff from this action is $E(v_i) - \alpha_i c_i$ since agent i will have to obtain information anyway. By requesting evidence first and then disregarding it, the principal and agent i get the same payoff.

In this case, it is easy to see that given the λ_i 's, the Lagrangian reduces to the Weitzman problem where the prize in box i is $v_i + \lambda_i$ and the cost of opening box i is $(\lambda_i + \alpha_i)c_i$. The incentive constraints are unaffected by this change, so the analysis is similar to the above.

6 Conclusion

There are several potentially interesting directions for future research. First, there are a number of simplifications here which could be relaxed. For example, we could consider situations where an agent might be unable to get evidence with some probability, as in the classic Dye (1985) model. Alternatively, we could consider situations where the principal prefers not to impose evidence gathering costs on the agents, all else equal. With this change, the no-free-lunch property is no longer guaranteed, so the appeal to Weitzman is, at the least, made more complex.

Second, our analysis focuses on the case where the principal has complete commitment power. Unlike in some work on mechanism design with evidence, this is *not* a setting where such an assumption is innocuous.⁸ In our mechanism, the principal sometimes gives the good to an agent even when he knows he would receive a higher payoff giving it to a different agent. Clearly, this cannot happen without commitment. One question of interest in the setting without commitment is whether the principal can show evidence from one agent to other agents in order to induce them to get evidence.

⁸For results on settings where commitment has no value to the principal, see, for example, Glazer and Rubinstein (2004, 2006), Hart, Kremer, and Perry (2017), or Ben-Porath, Dekel, and Lipman (2019, 2023).

Appendix

A Completion of Proof for Lemma 5

The completion of the proof of Lemma 5 follows from a result in Border (1991). For the reader's convenience, the Supplemental Appendix contains a complete proof of a version of Border's result. The version we use is the following. For this section of the Appendix and the related Supplemental Appendix only, we consider a different allocation problem with M agents where agent i has a *finite* set of types T_i . As in our model, the types are independently drawn across agents. Let μ_i denote the distribution over T_i . We consider allocations $P = (P_1, \ldots, P_N)$ with $P_i : T \to [0, 1]$ with $\sum_i P_i(t) \leq 1$ for all $t \in T$. Given P, we let $p = (p_1, \ldots, p_N)$ denote the *interim allocation* where

$$p_i(t_i) = \sum_{t_{-i} \in T_{-i}} \mu_{-i}(t_{-i}) P_i(t_i, t_{-i}).$$

When p and P are related in this fashion, we say P generates p.

A hierarchical allocation is an allocation P that can be constructed as follows. We have a ranking function R which maps $\bigcup_i T_i$ to $\{1, \ldots, K\}$ for some positive integer K. We assume that for every k < K, there is exactly one i such that $R(t_i) = k$ for some $t_i \in T_i$. Note that this restriction does not apply to rank K — there may be no or many agents with types at rank K.

Then given a type profile $t = (t_1, \ldots, t_M)$, either all agents have rank K or there is a unique i with $R(t_i) < R(t_j)$ for all $j \neq i$. If all agents have rank K, then $P_j(t) = 0$ for all j. If there is a unique i with $R(t_i) < R(t_j)$ for all $j \neq i$, then $P_i(t) = 1$. In other words, unless all agents are in the lowest rank, the agent who has the highest ranked type receives the good (where higher ranks have lower numbers).

We say that p is a *hierarchical interim allocation* if it is generated by a hierarchical allocation P. Of course, the collection of hierarchical interim allocations is a subset of the interim allocations. The following result is essentially Border's Lemma 6.1. We prove it in the Supplemental Appendix.

Theorem 5. Every interim allocation function p is a convex combination of hierarchical interim allocations.

We also use the following finite type version of Border's theorem which is the main result of Border (2007):

Theorem 6. p is an interim allocation function if and only if for every collection $\hat{T}_i \subseteq T_i$ for i = 1, ..., M, we have

$$\sum_{i} \sum_{t_i \in \hat{T}_i} p_i(t_i) \mu_i(t_i) \le 1 - \prod_{i} [1 - \mu_i(\hat{T}_i)].$$

We now use these results to complete the proof. So fix (e_1, \ldots, e_N) satisfying the Border conditions. We now show that there exist randomizations over the orderings generating these evidence probabilities.

First, consider tier 1, \mathcal{I}_1 . By assumption, the two Border conditions hold for tier 1, so the evidence probabilities for agents in this tier satisfy

$$\sum_{i \in \mathcal{I}_1} e_i [1 - F_i(\hat{v}_i)] = 1 - \prod_{i \in \mathcal{I}_1} F_i(\hat{v}_i)$$
(3)

and

$$\sum_{i\in\hat{\mathcal{I}}} e_i[1 - F_i(\hat{v}_i)] \le 1 - \prod_{i\in\hat{\mathcal{I}}} F_i(\hat{v}_i), \quad \forall \hat{\mathcal{I}} \subseteq \mathcal{I}_1.$$
(4)

Given this, consider the following alternative model. Let the set of agents be \mathcal{I}_1 . Each agent $i \in \mathcal{I}$ has two types, denoted ℓ_i and h_i , where $\mu_i(\ell_i) = F_i(\hat{v}_i)$. Define functions $\hat{p}_i: T_i \to [0, 1]$ for $i \in \mathcal{I}_1$ by

$$\hat{p}_i(t_i) = \begin{cases} 0, & \text{if } t_i = \ell_i; \\ e_i, & \text{if } t_i = h_i. \end{cases}$$

It is easy to see that equations (3) and (4) together with Theorem 6 imply that \hat{p} is an interim allocation function.

By Theorem 5, \hat{p} is a convex combination of hierchical interim allocations. That is, we can find q^1, \ldots, q^S and weights $\alpha_s \in (0, 1)$ with $\sum_s \alpha^s = 1$ with $\hat{p} = \sum_s \alpha^s q^s$. Clearly, the fact that $\hat{p}_i(\ell_i) = 0$ for all *i* implies that $q_i^s(\ell_i) = 0$ for all *i* and all *s*. By Theorem 6, we have

$$\sum_{i \in \mathcal{I}_1} q_i^s(h_i) [1 - F_i(\hat{v}_i)] \le 1 - \prod_{i \in \mathcal{I}_1} F_i(\hat{v}_i), \quad \forall s.$$

But

$$\sum_{s} \alpha_{s} \left\{ \sum_{i \in \mathcal{I}_{1}} q_{i}^{s}(h_{i}) [1 - F_{i}(\hat{v}_{i})] \right\} = \sum_{i \in \mathcal{I}_{1}} e_{i} [1 - F_{i}(\hat{v}_{i})] = 1 - \prod_{i \in \mathcal{I}_{1}} F_{i}(\hat{v}_{i}).$$

Hence

$$\sum_{i \in \mathcal{I}_1} q_i^s(h_i) [1 - F_i(\hat{v}_i)] = 1 - \prod_{i \in \mathcal{I}_1} F_i(\hat{v}_i), \quad \forall s$$

The left-hand side is the probability that one of the agents i is type h_i and receives the good, while the right-hand side is the probability that at least one agent is type h_i . Hence this equality says that for every s, q^s has the property that if at least one agent is type h_i , the good is allocated to one such agent.

Given this, consider any q^s . Since q^s is a hierarchical interim allocation, there is an allocation function, Q^s , and a ranking function, R^s , associated with it. From the above, we know that for any type profile such that some i is type h_i , the good is allocated to such an i. Because the allocation is hierarchical, there is a unique such i who gets the good with probability 1. Consider the allocation on type profile $(h_1, \ldots, h_{\#\mathcal{I}_1})$. Whichever agent i receives the good on this profile must have $R^s(h_i) = 1$. Let i_1 denote this agent.

Then consider the profile of types where agent i_1 is type ℓ_{i_1} and every other agent i is type h_i . Again, there must be an agent, say i_2 , who receives the good with probability 1 and hence we have $R^s(h_{i_2}) = 2$. Continuing in this fashion, we can construct the ranking R^s which completely ranks the h_i types of all agents. Define an ordering over $i \in \mathcal{I}_1, \succ^s$, by $i \prec^s j$ iff $R^s(h_i) < R^s(h_j)$.

By construction,

$$e_i = \hat{p}_i(h_i) = \sum_s \alpha^s q_i^s(h_i) = \sum_s \alpha^s \prod_{j \prec s_i} \mu_j(\ell_j) = \sum_s \alpha^s \prod_{j \prec s_i} F_j(\hat{v}_j).$$

Hence the randomization over orderings over \mathcal{I}_1 given by $O_1(\succ^s) = \alpha^s$ generates the e_i 's for \mathcal{I}_1 .

Next consider tier 2, \mathcal{I}_2 . By assumption, we know that the evidence probabilities for agents in this tier satisfy

$$\sum_{i \in \mathcal{I}_2} e_i [1 - F_i(\hat{v}_i)] = \left[\prod_{i \in \mathcal{I}_1} F_i(v_2^* - \lambda_i)\right] \left[1 - \prod_{i \in \mathcal{I}_1} F_i(\hat{v}_i)\right].$$
(5)

Also, we must have

$$\sum_{i\in\hat{\mathcal{I}}} e_i[1 - F_i(\hat{v}_i)] \le \left[\prod_{i\in\mathcal{I}_1} F_i(v_2^* - \lambda_i)\right] \left[1 - \prod_{i\in\hat{\mathcal{I}}} F_i(\hat{v}_i)\right], \quad \forall \hat{\mathcal{I}} \subseteq \mathcal{I}_2.$$
(6)

To see this, note that the left-hand side is the probability that an agent $i \in \hat{\mathcal{I}} \subseteq \mathcal{I}_2$ has $v_i \geq \hat{v}_i$ and receives the good, while the right-hand side is the probability that an agent $i \in \hat{\mathcal{I}}$ has $v_i \geq \hat{v}_i$ and that all the tier 1 agents have virtual values below v_2^* . Because an agent in tier 2 cannot get the good unless all tier 1 agents have virtual values below v_2^* , the left-hand side must be smaller than the right.

For $i \in \mathcal{I}_2$, let

$$\hat{e}_i = \frac{e_i}{\prod_{j \in \mathcal{I}_1} F_j(v_2^* - \lambda_i)}$$

Then equations (3) and (4) hold for tier 2 and the \hat{e} 's. That is, equations (5) and (6) can be rewritten as

$$\sum_{i \in \mathcal{I}_2} \hat{e}_i [1 - F_i(\hat{v}_i)] = 1 - \prod_{i \in \mathcal{I}_2} F_i(\hat{v}_i)$$

and

$$\sum_{i\in\hat{\mathcal{I}}} \hat{e}_i[1-F_i(\hat{v}_i)] \le 1 - \prod_{i\in\hat{\mathcal{I}}} F_i(\hat{v}_i), \quad \forall \hat{\mathcal{I}} \subseteq \mathcal{I}_2.$$

Hence the same argument as above shows that we can construct a probability distribution O_2 over orderings \succ^s over \mathcal{I}_2 such that for all $i \in \mathcal{I}_2$,

$$\hat{e}_i = \sum_s O_2(\succ^s) \prod_{j \prec^s i} F_j(\hat{v}_j)$$

or, equivalently,

$$e_i = \left[\prod_{j \in \mathcal{I}_1} F_j(v_2^* - \lambda_i)\right] \sum_s O_2(\succ^s) \prod_{j \prec^s i} F_j(\hat{v}_j).$$

Iterating this argument for the remaining tiers completes the proof.

B Proof of Theorem 3

The proof is by contradiction. So suppose $c_i = c_j$ and $F_i = F_j$, but we have an optimal mechanism d^1 with $e_i(d^1) = e_i^1 = 0$ and $e_j(d^1) = e_j^1 > 0$. By the no-free-lunch property, we must have $P_i(v \mid d^1) = 0$ for (almost) all v. Given this, it's not hard to see we must have at least three agents. Otherwise, the best outcome of this form is to simply give the good to j with probability 1. Essentially the same argument as the proof of Lemma 1 gives a contradiction to this being optimal.

We write \hat{v}_k^1 , λ_k^1 , etc., to denote the relevant variables for this outcome. We write P^1 for $P(d^1)$ and e^1 for $e(d^1)$.

Let d^2 denote the mechanism which flips the roles of *i* and *j*, so *j* is never asked for evidence and *i* is asked in the situations in which *j* had been in d^1 . Let $d^* = (1/2)d^1 + (1/2)d^2$, the 50–50 randomization between these mechanisms. We write P^* for $P(d^*)$, etc. Key observation: $e_i^* = e_j^* = (1/2)e_j^1 > 0$. This means that we must have $\hat{v}_i^* + \lambda_i^* = \hat{v}_j^* + \lambda_j^*$. Because these are defined from the same function which is strictly decreasing in λ in the relevant range, this implies $\lambda_i^* = \lambda_j^*$ and $\hat{v}_i^* = \hat{v}_j^*$.

So what is the mechanism d^* ? This mechanism has a certain set of histories on which either *i* or *j* is next asked for evidence, each with probability 1/2. (By history here, we include the randomization over the order if *i* and *j* are in a tier with one or more other agents.) But if, say, *i* is chosen, the mechanism *never* continues to *j*. Because we only get to *i* if all previously observed virtual values are below $\hat{v}_i^* + \lambda_i^* = \hat{v}_j^* + \lambda_j^*$, the only way this can be true is if the good is given to *i* with probability 1 — that is, if $\hat{v}_i^* \leq 0$. But this implies *i* receives the good iff she is asked for evidence. Hence her expected utility in the mechanism is $e_i^* - c_i e_i^* = (1 - c_i) e_i^*$. We know that $e_i^* > 0$ and $c_i < 1$, so *i*'s expected payoff is strictly positive, implying that her incentive constraint is not binding. Hence $\lambda_i^* = 0$, so \hat{v}_i^* satisfies

$$0 = \int_{\hat{v}_i^*}^1 (v - \hat{v}_i^*) f_i(v) \, dv$$

requiring $\hat{v}_i^* = 1$, a contradiction.

C Proof of Theorem 4

The following lemmas will be useful.

Recall that \hat{v}_i is defined from λ_i by

$$\lambda_{i} = \frac{1}{c_{i}} \int_{\hat{v}_{i}}^{1} (v - \hat{v}_{i}) f_{i}(v) \, dv \tag{7}$$

It is more convenient to work with the inverse function, defining λ_i from \hat{v}_i .

Lemma 6. Define a function $\lambda_n(\cdot)$ by

$$\lambda_n(\hat{v}) \equiv \frac{1}{c_n} \int_{\hat{v}}^1 (v - \hat{v}) f_n(v) \, dv$$

Then $\lambda_n(\cdot)$ is decreasing. Also, the function $\hat{v} + \lambda_n(\hat{v})$ is increasing in \hat{v} , weakly at $\hat{v} = \hat{v}_n$, strictly for $\hat{v} > \hat{v}_n$, where \hat{v}_n is the value at the optimal mechanism.

Proof. It's easy to see that

$$\lambda_n'(\hat{v}) = -\frac{1 - F_n(\hat{v})}{c_n} \le 0$$

Also,

$$\frac{\partial}{\partial \hat{v}} \left[\hat{v} + \lambda_n(\hat{v}) \right] = 1 - \frac{1 - F_n(\hat{v})}{c_n}.$$

This is weakly positive iff $c_n \ge 1 - F_n(\hat{v})$. We now show that this must hold for all $\hat{v} \ge \hat{v}_n$, strictly for $\hat{v} > \hat{v}_n$.

First, suppose $\lambda_n = 0$. Then $\hat{v}_n = 1$ since this is the only way to solve equation (7). Hence $F_n(\hat{v}_n) = 1$ so

$$\frac{\partial}{\partial \hat{v}} \left[\hat{v} + \lambda_n(\hat{v}_n) \right] = 1 > 0.$$

Second, suppose $\lambda_n > 0$. Then *n*'s expected payoff in the optimal mechanism is zero. Hence we must have

$$e_n[1 - F_n(\hat{v}_n) - c_n] \le 0$$

as agent *n* certainly gets the good if asked for evidence when her value is above \hat{v}_n and *n* might get the good even when her value is below \hat{v}_n . Hence $1 - F_n(\hat{v}_n) \leq c_n$, implying the derivative above is weakly positive at $\hat{v} = \hat{v}_n$. For $\hat{v} > \hat{v}_n$, we have

$$1 - F_n(\hat{v}) < 1 - F_n(\hat{v}_n) \le c_n,$$

so, as claimed, the derivative is positive at $\hat{v} > \hat{v}_n$.

Lemma 7. For all i and v_i , $p_i(v_i \mid d) \leq e_i$.

Proof. In the proof of Lemma 4, we showed that

$$p_i(v_i \mid d) = \begin{cases} e_i, & \text{if } v_i \ge \hat{v}_i; \\ \prod_{j \ne i \mid \hat{v}_j + \lambda_j \ge v_i + \lambda_i} F_j(v_i + \lambda_i - \lambda_j), & \text{otherwise}, \end{cases}$$

Hence the claim of the lemma is obviously true for $v_i \ge \hat{v}_i$. For $v_i < \hat{v}_i$, it is clear that $p_i(v_i \mid d)$ is increasing, so

$$p_i(v_i \mid d) \le \lim_{v'_i \uparrow \hat{v}_i} p_i(v'_i \mid d).$$

Clearly, the limit on the right-hand side is the probability that every agent j in the same tier as i or a higher tier has a virtual value below $\hat{v}_i + \lambda_i$. In this event, the optimal mechanism *must* ask i for evidence and may ask i for evidence even before this. Hence this probability must be weakly smaller than e_i .

Lemma 8. If i's incentive constraint is not binding, then $\lambda_i = 0$, $\hat{v}_i = 1$, and $e_i = 1$. If i's incentive constraint is binding, then $\hat{v}_i + \lambda_i < 1$.

Proof. Obviously, if *i*'s incentive constraint doesn't bind, the Lagrange multiplier for that constraint, λ_i , is 0. Recall that \hat{v}_i is defined by

$$\lambda_i c_i = \int_{\hat{v}_i}^1 (v - \hat{v}_i) f_i(v) \, dv,$$

so if $\lambda_i = 0$, we must have $\hat{v}_i = 1$. Hence *i*'s index, $\hat{v}_i + \lambda_i$, equals 1.

So consider any $j \neq i$. Then j's index satisfies

$$\hat{v}_j + \lambda_j = \hat{v}_j + \frac{1}{c_j} \int_{\hat{v}_j}^1 (v - \hat{v}_j) f_j(v) \, dv.$$

By Lemma 6, viewing the right-hand side as a function of \hat{v}_j , it is increasing in \hat{v}_j at the value in the optimal mechanism. Hence

$$\hat{v}_j + \frac{1}{c_j} \int_{\hat{v}_j}^1 (v - \hat{v}_j) f_j(v) \, dv \le 1 + \frac{1}{c_j} \int_1^1 (v - 1) f_j(v) \, dv = 1,$$

with a strict inequality if $\hat{v}_j < 1$. Hence *i* must be in the highest tier and any agent $j \neq i$ in the same tier must also have $\lambda_j = 0$ and $\hat{v}_j = 1$. Therefore, with probability 1, all of the agents in this highest tier will be asked for evidence, so $e_i = 1$.

Finally, if *i*'s incentive constraint is binding, we have $\lambda_i > 0$ and hence $\hat{v}_i < 1$. From the argument regarding $j \neq i$ in the previous paragraph, we see that this implies $\hat{v}_i + \lambda_i < 1$.

Recall that *i* is stronger than *j* if $c_i \leq c_j$ and F_i FOSD F_j . The definition implies that if *i* is stronger than *j* and *j* is stronger than *i*, then we must have $c_i = c_j$ and $F_i = F_j$.

The next lemma shows the result of Theorem 4 for a special case.

Lemma 9. If d is an optimal mechanism, i is stronger than j, and $e_i(d) = 0$, then we have $e_i(d) = 0$.

Proof. Suppose not. So suppose *i* is stronger than *j*, $e_i = 0$, and $e_j > 0$. (We omit the *d* argument when no confusion is likely to result.) By Theorem 3, we cannot have $c_i = c_j$ and $F_i = F_j$, so either $c_i < c_j$ or F_i strictly FOSD F_j or both.

By the no-free-lunch property, $e_i = 0$ implies that *i* never receives the good.

We construct an alternative mechanism \bar{d} as follows. Define a function $\varphi: V_i \to V_j$ by

$$\varphi(v_i) = F_j^{-1}(F_i(v_i)).$$

Define \bar{d} to be exactly the same as d except for two changes. First, on any history where d asks j for evidence with positive probability, \bar{d} asks i instead with this same probability. Second, in the event that i is asked for evidence and proves value v_i , the mechanism \bar{d} treats this the same way mechanism d would treat proof by j of value $\varphi(v_i)$.

The key observation is that the distribution of $\varphi(v_i)$ is the same as the distribution of v_i . That is, for any z, we have

$$\Pr[\varphi(v_i) \le z] = \Pr[F_j^{-1}(F_i(v_i)) \le z]$$

=
$$\Pr[F_i(v_i) \le F_j(z)]$$

=
$$\Pr[v_i \le F_i^{-1}(F_j(z))]$$

=
$$F_i(F_i^{-1}(F_j(z))) = F_j(z)$$

Hence for any $k \neq i, j$, the probability k is asked for evidence and the probability k receives the good is unaffected.

Similarly, the probability i gets the good in \overline{d} is the same as the probability j got the good in d and the probability i is asked for evidence in \overline{d} is the same as the probability j is asked in d. Because $c_i \leq c_j$, then, i's incentive constraint is satisfied. (The incentive constraint for j is trivially satisfied in \overline{d} as she is never asked for evidence.)

In short, \overline{d} is incentive compatible.

Finally, consider the principal's payoff from \overline{d} . Because the probability the principal gives the good to any agent $k \neq i, j$ is unchanged, we can write the payoff as

$$\mathbf{E}_{v}[P_{i}(v_{i}, v_{j}, v_{-ij} \mid \bar{d})v_{i}] + \sum_{k \neq i,j} \mathbf{E}_{v}[P_{k}(v \mid d)v_{k}].$$

We can write the first term as

$$\mathbf{E}_{v}[P_{j}(v_{j},\varphi(v_{i}),v_{-ij}\mid d)v_{i}].$$

Since $\varphi(v_i)$ has the same distribution as v_j , we can write the corresponding term in the original mechanism as

$$\mathbf{E}_{v}[P_{j}(v_{i}, v_{j}, v_{-ij} \mid d)v_{j}] = \mathbf{E}_{v}[P_{j}(v_{j}, \varphi(v_{i}), v_{-ij} \mid d)\varphi(v_{i})].$$

Hence the new mechanism yields the principal a weakly better payoff if

$$v_i \ge \varphi(v_i), \quad \forall v_i$$

or $v_i \ge F_j^{-1}(F_i(v_i))$ or $F_j(v) \ge F_i(v)$ for all v. This holds as F_i FOSD F_j .

Hence \overline{d} yields at least as high a payoff for the principal as d. If $F_i \neq F_j$, the gain must be strict.

So suppose $F_i = F_j$. Then we must have $c_i < c_j$ by Theorem 3. Then the incentive constraint for *i* is not binding in (\bar{P}, \bar{e}) . If *j*'s incentive constraint in (P, e) was binding, then again it must be possible to make the principal strictly better off than at *d*. This is because we have effectively replaced *j* with an equivalent agent and turned a binding constraint into a nonbinding one, ensuring that an improvement is possible. Hence if *d* was optimal, *j*'s incentive constraint was not binding.

In this case, i and j are effectively identical from the point of view of the principal. Even though $c_i < c_j$, the fact that j's incentive constraint does not bind implies that reductions in j's cost have no effect on the optimal mechanism. Hence we can analyze this case as if $c_i = c_j$. By Theorem 3, then, we have the needed contradiction.

Lemma 10. If *i* is stronger than *j* and $0 < e_i < e_j$, then $\hat{v}_j + \lambda_j \ge \hat{v}_i + \lambda_i$, $\hat{v}_j \ge \hat{v}_i$, and $\lambda_j \le \lambda_i$.

Proof of Lemma. Clearly, if $e_j > e_i$, we cannot have j in a lower tier than i. Any agent is only asked for evidence after all agents in higher tiers are asked, so $e_j > e_i$ implies j is in a weakly higher tier than i. Hence

$$\hat{v}_j + \lambda_j \ge \hat{v}_i + \lambda_i,$$

establishing the first claim.

Next, we show that if i is stronger than j, then for all \hat{v} , we have

$$\lambda_i(\hat{v}) \ge \lambda_i(\hat{v}),$$

where for any n, $\lambda_n(\cdot)$ is the function defined in Lemma 6. To see this, first note that this holds with equality if $c_i = c_j$ and $F_i = F_j$ since the functions are then the same. So suppose $F_i = F_j$ but $c_i < c_j$. Then the λ_i function must be larger at every $\hat{v} < 1$ since it's defined by dividing by a strictly smaller number. Alternatively, suppose $c_i = c_j$ but F_i FOSD F_j . Because the function max $\{0, v - \hat{v}\}$ is increasing in v, F_i FOSD F_j implies

$$\int_{\hat{v}}^{1} (v - \hat{v}) f_i(v) \, dv \ge \int_{\hat{v}}^{1} (v - \hat{v}) f_j(v) \, dv,$$

again implying $\lambda_i(\hat{v}) \geq \lambda_j(\hat{v})$.

Suppose, contrary to what we claim, that $\hat{v}_i > \hat{v}_j$. Then by Lemma 6, we know that $\hat{v} + \lambda_j(\hat{v})$ is strictly increasing in \hat{v} for all $\hat{v} \in (\hat{v}_j, \hat{v}_i)$. Hence we must have

$$\hat{v}_j + \lambda_j(\hat{v}_j) < \hat{v}_i + \lambda_j(\hat{v}_i).$$

But i stronger than j then implies

$$\hat{v}_j + \lambda_j(\hat{v}_j) < \hat{v}_i + \lambda_j(\hat{v}_i) \le \hat{v}_i + \lambda_i(\hat{v}_i).$$

From equation (7), we have

$$\hat{v}_j + \lambda_j = \hat{v}_j + \lambda_j(\hat{v}_j) < \hat{v}_i + \lambda_j(\hat{v}_i) \le \hat{v}_i + \lambda_i(\hat{v}_i) = \hat{v}_i + \lambda_i.$$

But this contradicts j being in a weakly higher tier than i. Hence $\hat{v}_j \geq \hat{v}_i$, establishing the second claim.

Finally, because the $\lambda_n(\cdot)$ functions are decreasing, equation (7) implies

$$\lambda_i = \lambda_i(\hat{v}_i) \ge \lambda_i(\hat{v}_j).$$

Again using i stronger than j, we have

$$\lambda_i(\hat{v}_j) \ge \lambda_j(\hat{v}_j) = \lambda_j.$$

Hence $\lambda_i \geq \lambda_j$, establishing the third claim.

We prove Theorem 4 by contradiction. So suppose *i* is stronger than *j* but $e_i < e_j$. By Lemma 9, we must have $e_i > 0$. By Lemma 8, the fact that $e_i < e_j \leq 1$ implies that *i*'s incentive constraint is binding. By Lemma 8, this implies $\hat{v}_i + \lambda_i < 1$. Also, by Lemma 10, we have $\hat{v}_j + \lambda_j \geq \hat{v}_i + \lambda_i$, $\hat{v}_j \geq \hat{v}_i$, and $\lambda_j \leq \lambda_i$.

The proof of Lemma 4 shows that

$$p_i(v_i \mid d) = \begin{cases} e_i, & \text{if } v_i \ge \hat{v}_i; \\ \prod_{j \ne i \mid \hat{v}_j + \lambda_j \ge v_i + \lambda_i} F_j(v_i + \lambda_i - \lambda_j), & \text{otherwise,} \end{cases}$$

so we can write the expected payoff to any agent k as

$$U_k = e_k [1 - F_k(\hat{v}_k)] - e_k c_k + \int_0^{\hat{v}_k} \left[\prod_{j \neq k \mid \hat{v}_j + \lambda_j \ge v_k + \lambda_k} F_j(v + \lambda_k - \lambda_j) \right] f_k(v) \, dv.$$

Because *i*'s incentive constraint binds, we must have $U_i = 0 \le U_j$.

For $v_j < \hat{v}_j$, we have

$$p_j(v_j \mid d) = \prod_{k \neq j \mid \hat{v}_k + \lambda_k \ge v_j + \lambda_j} F_k(v_j + \lambda_j - \lambda_k).$$

Let $w = v_j + \lambda_j$. Then for $w \in [\lambda_i, \hat{v}_j + \lambda_j)$, we have

$$p_{j}(w - \lambda_{j} \mid d) = \prod_{\substack{k \neq j \mid \hat{v}_{k} + \lambda_{k} \ge w}} F_{k}(w - \lambda_{k})$$

$$= F_{i}(w - \lambda_{i}) \prod_{\substack{k \neq i, j \mid \hat{v}_{k} + \lambda_{k} \ge w}} F_{k}(w - \lambda_{k})$$

$$\leq F_{j}(w - \lambda_{j}) \prod_{\substack{k \neq i, j \mid \hat{v}_{k} + \lambda_{k} \ge w}} F_{k}(w - \lambda_{k})$$

$$= \prod_{\substack{k \neq i \mid \hat{v}_{k} + \lambda_{k} \ge w}} F_{k}(w - \lambda_{k})$$

$$= p_{i}(w - \lambda_{i} \mid d), \quad \forall w - \lambda_{i} < \hat{v}_{i}$$

where the inequality in the third line comes from F_i FOSD F_j and $\lambda_i \geq \lambda_j$. To understand the last line, note that $w - \lambda_i < \hat{v}_i$ means that $v_i = w_i - \lambda_i < \hat{v}_i$, so the formula given in the proof of Lemma 4 for $p_i(v_i \mid d)$ applies.

Summarizing, we have

$$p_j(v_j \mid d) \le p_i(v_j + \lambda_j - \lambda_i \mid d), \quad \forall v_j < \hat{v}_i + \lambda_i - \lambda_j.$$

Hence, by Lemma 7, we have

$$p_j(v_j \mid d) \le e_i, \quad \forall v_j < \hat{v}_i + \lambda_i - \lambda_j \le \hat{v}_j.$$

We derive a further implication shortly. First, because $p_i(v_i \mid d)$ is increasing in v_i , the hypothesis that F_i FOSD F_j implies

$$\int_0^1 p_i(v_i \mid d) f_i(v_i) \, dv_i \ge \int_0^1 p_i(v_i \mid d) f_j(v_i) \, dv_i.$$

Rewriting the right-hand side, we have

$$\int_{0}^{1} p_{i}(v_{i} \mid d) f_{j}(v_{i}) dv_{i} = \int_{0}^{\hat{v}_{i}} \left[\prod_{k \neq i \mid \hat{v}_{k} + \lambda_{k} \ge v_{i} + \lambda_{i}} F_{k}(v_{i} + \lambda_{i} - \lambda_{k}) \right] f_{j}(v_{i}) dv_{i} + [1 - F_{j}(\hat{v}_{i})]e_{i}.$$

Change variables in the integral on the right-side side by defining $w = v_i + \lambda_i$. This changes the integral to

$$\int_{\lambda_i}^{\hat{v}_i + \lambda_i} \left[\prod_{k \neq i \mid \hat{v}_k + \lambda_k \ge w} F_k(w - \lambda_k) \right] f_j(w - \lambda_i) \, dw.$$

The same reasoning as above shows that this is weakly larger than

$$\int_{\lambda_i}^{\hat{v}_i + \lambda_i} \left[\prod_{k \neq j \mid \hat{v}_k + \lambda_k \ge w} F_k(w - \lambda_k) \right] f_j(w - \lambda_i) \, dw.$$

Change variables in the integral again, replacing w with $v_j + \lambda_i$ to write this as

$$\int_0^{\hat{v}_i} \left[\prod_{k \neq j \mid \hat{v}_k + \lambda_k \ge v_j + \lambda_i} F_k(v_j + \lambda_i - \lambda_k) \right] f_j(v_j) \, dv_j.$$

The fact that $\lambda_i \geq \lambda_j$ implies

$$F_k(v_j + \lambda_i - \lambda_k) \ge F_k(v_j + \lambda_j - \lambda_k)$$

for all k and v_j . Also, note that if we change the index on the product from $k \neq j$ such that $\hat{v}_k + \lambda_k \geq v_j + \lambda_i$ to $k \neq j$ such that $\hat{v}_k + \lambda_k \geq v_j + \lambda_j$, the fact that $\lambda_i \geq \lambda_j$ means that we will be taking the product over weakly more k's. Since each term in the product is weakly less than 1, this must reduce the product. Hence the expression above is weakly larger than

$$\int_0^{\hat{v}_i} \left[\prod_{k \neq j \mid \hat{v}_k + \lambda_k \ge v_j + \lambda_j} F_k(v_j + \lambda_j - \lambda_k) \right] f_j(v_j) \, dv_j = \int_0^{\hat{v}_i} p_j(v_j \mid d) f_j(v_j) \, dv_j.$$

The fact that i's incentive constraint binds implies

$$c_i e_i = \int_0^1 p_i(v_i \mid d) f_i(v_i) \, dv_i \ge \int_0^{\hat{v}_i} p_j(v_j \mid d) f_j(v_j) \, dv_j + [1 - F_j(\hat{v}_i)] e_i.$$

j's incentive constraint implies

$$\int_0^1 p_j(v_j \mid d) f_j(v_j) \, dv_j + (1 - F_j(\hat{v}_j)) e_j \ge c_j e_j.$$

For $v_j \in [\hat{v}_i, \hat{v}_i + \lambda_i - \lambda_j]$, we know that $p_j(v_j \mid d) \leq e_i$. For $v_j \in [\hat{v}_i + \lambda_i - \lambda_j, \hat{v}_j]$, we have $p_j(v_j \mid d) \leq e_j$. Hence

$$\int_{0}^{\hat{v}_{i}} p_{j}(v_{j} \mid d) f_{j}(v_{j}) \, dv_{j} + [F_{j}(\hat{v}_{i} + \lambda_{i} - \lambda_{j}) - F_{j}(\hat{v}_{i})]e_{i} + [1 - F_{j}(\hat{v}_{i} + \lambda_{i} - \lambda_{j})]e_{j} \ge c_{j}e_{j}.$$

Summarizing, we have

$$\begin{split} c_i e_i &\geq \int_0^{\hat{v}_i} p_j(v_j \mid d) f_j(v_j) \, dv_j + [1 - F_j(\hat{v}_i)] e_i \\ &\geq [1 - F_j(\hat{v}_i)] e_i + c_j e_j - [F_j(\hat{v}_i + \lambda_i - \lambda_j) - F_j(\hat{v}_i)] e_i - [1 - F_j(\hat{v}_i + \lambda_i - \lambda_j)] e_j \\ &= c_j e_j + [1 - F_j(\hat{v}_i + \lambda_i - \lambda_j)] e_i - [1 - F_j(\hat{v}_i + \lambda_i - \lambda_j)] e_j. \end{split}$$

$$[c_i - (1 - F_j(\hat{v}_i + \lambda_i - \lambda_j))]e_i \ge [c_j - (1 - F_j(\hat{v}_i + \lambda_i - \lambda_j))]e_j.$$
(8)

Recall that

$$U_{i} = 0 = (1 - F_{i}(\hat{v}_{i}) - c_{i})e_{i} + \int_{0}^{\hat{v}_{i}} \left[\prod_{k \neq i | \hat{v}_{k} + \lambda_{k} \ge v_{i} + \lambda_{i}} F_{k}(v_{i} + \lambda_{i} - \lambda_{k})\right] f_{i}(v_{i}) dv_{i}.$$

Obviously, the integral is non-negative as it is a probability. We now show that $e_i > 0$ implies that the integral must be strictly positive. Recall that the mechanism does not ask *i* for evidence until every agent *k* with $\hat{v}_k + \lambda_k > \hat{v}_i + \lambda_i$ has already been asked for evidence and has been found to have a virtual value strictly below $\hat{v}_i + \lambda_i$. Hence the fact that $e_i > 0$ implies that this event must have positive probability.

Furthermore, any other agent k in the same tier as i must have a positive probability of having a virtual value below the tier threshold. To see this, suppose it were not true. k's virtual value is below the threshold for the tier if $v_k < \hat{v}_k$, so the only way this could have zero probability is if $\hat{v}_k = 0$. But if $\hat{v}_k = 0$, this means agent k receives the good if and only if she is asked for evidence, so her probability of getting the good ex ante is e_k . But then her expected payoff is $e_k(1 - c_k) > 0$ as we assume $c_k < 1$ for all k. This implies that k's incentive constraint is not binding, which by Lemma 8 implies $\hat{v}_k = 1$, not zero, a contradiction.

Hence $U_i = 0$ implies

$$0 < \int_0^{\hat{v}_i} \left[\prod_{k \neq i \mid \hat{v}_k + \lambda_k \ge v_i + \lambda_i} F_k(v_i + \lambda_i - \lambda_k) \right] f_i(v_i) \, dv_i = [c_i - (1 - F_i(\hat{v}_i))] e_i$$

So we have

$$c_i > 1 - F_i(\hat{v}_i) \ge 1 - F_i(\hat{v}_i + \lambda_i - \lambda_j) \ge 1 - F_j(\hat{v}_i + \lambda_i - \lambda_j),$$

where the second inequality is implies by $\lambda_i \geq \lambda_j$ and the second from F_i FOSD F_j .

Hence $c_i - [1 - F_j(\hat{v}_i + \lambda_i - \lambda_j)] > 0$, so $e_j > e_i$ implies

$$[c_i - (1 - F_j(\hat{v}_i + \lambda_i - \lambda_j))]e_i < [c_i - (1 - F_j(\hat{v}_i + \lambda_i - \lambda_j))]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j \le [c_j - (1 - F_j(\hat{v}_i + \lambda_j - \lambda_j)]e_j$$

where the second inequality follows from $c_j \ge c_i$. But this contradicts equation (8).

So

References

- [1] Ball, I., and D. Kattwinkel, "Probabilistic Verification in Mechanism Design," working paper, current draft May 2022.
- [2] Ben-Porath, E., E. Dekel, and B. Lipman, "Mechanisms with Evidence: Commitment and Robustness," *Econometrica*, 87, March 2019, 529–566.
- [3] Ben-Porath, E., E. Dekel, and B. Lipman, "Mechanism Design for Acquisition of/Stochastic Evidence," working paper, current draft June 2023.
- [4] Bergemann, D., and J. Välimäki, "Information in Mechanism Design," in R. Blundell, W. Newey, and T. Persson, eds., Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press, 2006, 186–221.
- [5] Border, K., "Implementation of Reduced Form Auctions: A Geometric Approach," *Econometrica*, 59, July 1991, 1175–1187.
- [6] Border, K., "Reduced Form Auctions Revisited," *Economic Theory*, **31**, April 2007, 167–181.
- [7] Bull, J., and J. Watson, "Hard Evidence and Mechanism Design," Games and Economic Behavior, 58, January 2007, 75–93.
- [8] Che, Y.-K., and N. Kartik, "Opinions as Incentives," *Journal of Political Economy*, 117, October 2009, 815–860.
- [9] Che, Y.-K., J. Kim, and K. Mierendorff, "Generalized Reduced-Form Auctions: A Network-Flow Approach," *Econometrica*, 81, November 2013, 2487–2520.
- [10] Crémer, J., Y. Spiegel, and C. Zheng, "Auctions with Costly Information Acquisition," *Economic Theory*, 38, January 2009, 41–72.
- [11] DeMarzo, P., I. Kremer, and A. Skrzypacz, "Test Design and Minimum Standards," *American Economic Review*, 109, June 2019, 2173–2207.
- [12] Deneckere, R. and S. Severinov, "Mechanism Design with Partial State Verifiability," Games and Economic Behavior, 64, November 2008, 487–513.
- [13] Dye, R. A., "Disclosure of Nonproprietary Information," Journal of Accounting Research, 23, 1985, 123–145.
- [14] Gershkov, A., and B. Szentes, "Optimal Voting Schemes with Costly Information Acquisition," *Journal of Economic Theory*, 144, January 2009, 36–68.

- [15] Glazer, J., and A. Rubinstein, "On Optimal Rules of Persuasion," *Econometrica*, 72, November 2004, 1715–1736.
- [16] Glazer, J., and A. Rubinstein, "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach," *Theoretical Economics*, 1, December 2006, 395–410.
- [17] Green, J., and J.-J. Laffont, "Partially Verifiable Information and Mechanism Design," *Review of Economic Studies*, 53, July 1986, 447–456.
- [18] Hart, S., I. Kremer, and M. Perry, "Evidence Games: Truth and Commitment," American Economic Review, 107, March 2017, 690–713.
- [19] Mierendorff, K., "Asymmetric Reduced Form Auctions," *Economics Letters*, 110, January 2011, 41–44.
- [20] Shishkin, D., "Evidence Acquisition and Voluntary Disclosure," working paper, December 2020.
- [21] Weitzman, M., "Optimal Search for the Best Alternative," Econometrica, 47, May 1979, 641–654.