European University Institute

MAX WEBER PROGRAMME FOR POSTDOCTORAL STUDIES

EUI WORKING PAPERS

MWP 2017/10
Max Weber Programme

Hard Evidence and Welfare in Adverse Selection Environments

Kym Pram

European University Institute
Max Weber Programme

Hard Evidence and Welfare in Adverse Selection Environments

Kym Pram

EUI Working Paper **MWP** 2017/10

ISSN 1830-7728

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Abstract

I consider environments in which an agent with private information can acquire arbitrary hard evidence about his type before interacting with a principal. In a broad class of screening models, I show that there is always an evidence structure that interim Pareto improves over the no-evidence benchmark whenever some types of the agent take an outside option in the benchmark case, and additional weak conditions, including either a single-crossing condition or state-independence of the principal's payoffs, are satisfied. I show that the sufficient conditions are tight and broadly applicable. Addressing concerns about multiple equilibria, I show how a planner can restrict the available evidence to ensure that an equilibrium which interim Pareto-improves over the benchmark case is obtained. Furthermore, I show that Pareto-improving evidence can arise endogenously when agents choose what evidence to acquire (and disclose).

Keywords

Information Economics; Hard Evidence; Mechanism Design; Microeconomic Theory; Insurance.

I would like to thank Eddie Dekel, Jeff Ely and Alessandro Pavan for their invaluable support and advice. I would also like to thank Navin Kartik, Francesco Giovannoni, Anne-Katrin Roessler, Maryam Saeedi, Piero Gottardi, Andrea Galeotti, Chris Li and seminar participants at Northwestern University, Bocconi University, the EUI, Oxford University and the University of Bristol for helpful comments and suggestions.

Kym Pram
Max Weber Programme, 2016-2017
European University Institute. Email: kym.pram@eui.eu

1 Introduction

In 2008 the US Congress passed the Genetic Information Nondiscrimination Act (GINA). Among other provisions, the GINA prohibits health insurers from denying coverage or charging higher premiums based on a genetic predisposition to develop a disease in the future. The issue had become prominent following technological advances in genetic testing, which, while offering only crude predictions for a limited set of conditions at present, promises to become more accurate and broadly applicable in the future. Genetic testing is a form of *evidence*: information which is true, or at least costly to falsify (unlike cheap talk) and which must be voluntarily disclosed (unlike directly observable information).[1]

New technologies such as genetic testing will expand the scope and accuracy of evidence used in economic transactions, but evidence has long been ubiquitous in economic life. Student ID cards used to obtain discounts at a movie theater, academic qualifications presented by a job seeker to a potential employer, and quotes from suppliers shared by a contractor are all examples. In general, evidence can ameliorate inefficiencies due to asymmetric information. However, this may be at consumers' expense: for example, a concern driving the GINA may be that evidence about genetic predispositions to costly diseases would lead to unaffordably high premiums for some consumers. In this paper, I investigate the welfare implications of (endogenous) evidence and the optimal availability of evidence for consumers.

This paper considers evidence, which is distinct from two types of information more commonly studied in economic models: cheap talk and observable information. Claims backed by evidence must be true, or at least bear some relation to the truth, unlike cheap talk which acquires a relation to the truth only endogenously. For example, a firm can claim that it can purchase some input at a certain price regardless of whether or not it can (cheap talk), but can only provide a quote from a supplier if it can purchase at that price (evidence). In contrast to observable information, the holder of evidence chooses voluntarily to disclose it. For example, a movie theater cannot tell whether a customer is a student unless a student ID is presented, but may be able to observe directly whether a customer is a child.

My model applies to a broad class of environments in which an informed agent can provide evidence of his type to a principal. I show that, under weak conditions, there is an evidence structure that interim Pareto-improves on the equilibrium when no evidence is available. The key property is that the agent has access to an outside option that is taken by some types in an optimal mechanism without evidence (some other broadly applicable monotonicity and single-crossing conditions are required in addition).

Although the requirement that some types take the outside option without evidence is not a primitive condition, it is easy to check in the many settings where optimal mechanisms without evidence are well understood. Examples of environments that satisfy the conditions include general buyer seller interactions (generalizing from the unit-valuation-zero-cost environment studied in the leading case to allow for sales of multiple goods and any type-independent cost function), and monopoly insurance.

[1] In reality, consumers may also *learn* about their types by obtaining a genetic test. This aspect is not modeled in this paper. Nevertheless, the basic mechanism behind my results will continue to apply as long as consumers have some private information pre-testing. Moreover, adding learning to the model could be expected to strengthen the favorable welfare results I obtain for two reasons: Firstly, since learning implies that evidence-acquisitions are made at an ex-ante stage, the consumer can make an ex-ante optimal tradeoff between the welfare effects on possible interim types. Secondly, learning per-se is typically welfare improving.

In a motivating example, I show that a buyer facing a monopolistic seller can be made better off when evidence is available compared to the benchmark case when no evidence is available. In addition, the same evidence structure can arise endogenously when the buyer chooses what evidence to acquire before the monopolist sets a price, and this result is robust to several different assumptions about the timing of evidence acquisition and presentation. Modelling evidence acquisition as endogenous is appropriate in many real-life cases: for example, students can decide whether to acquire a student ID, firms have some flexibility over the accounting information they collect, patients choose whether to undergo genetic testing.

Surprisingly, although the monopolist could extract the full consumer surplus if the buyer acquired all available evidence, there is always an equilibrium in which the buyer acquires partial evidence and is made better off. Moreover, if evidence is made available by a planner with consumer surplus as her objective – for example student IDs provided by a university or senior cards provided by a government – the planner can choose an information structure which simultaneously allows for an equilibrium that interim Pareto dominates the equilibrium with no evidence and rules out equilibria that make any type of the consumer worse off. Under a mild refinement – iterated admissibility – the planner can induce the 'good equilibrium' as the unique outcome.

The intuition in the buyer-monopolist case can be illustrated using a stylized example. A movie theater serves both students, who are willing to pay either eight or nine dollars for a ticket and non-students, who are willing to pay either ten or eleven dollars. The proportions of willingness-to-pay types are such that the revenue-maximizing price when the monopolist is constrained to charge the same price to all buyers is ten dollars. At that price students do not go to the movies.

Now suppose that the students acquire ID cards and the theater is aware of this. The theater now charges eight dollars for customers presenting a student ID card and continues to charge ten dollars for customers who do not present an ID card. Students are strictly better off and non-students are no worse off. At the same time, if it were known that every customer had acquired evidence that fully identified their willingness to pay, the theater could engage in first-degree price discrimination by requiring the customer to disclose his willingness to pay in order to make any purchase.

A naive analysis may suggest that a student discount implies a 'non-student surcharge': if students are willing to pay less than the average then non-students are willing to pay more, and so will be charged more if students identify themselves. In fact, as long as students are excluded when they cannot be identified, this is incorrect. If the seller would prefer to charge a higher price to non-students when they do not present a student ID it would also have been feasible to do so, and would have had the same effects on revenue, when the student ID did not exist.

More formally, when a student ID can be presented, incentive constraints between the student buyers and the non-student buyers can be ignored. However, in the benchmark case, since the students receive the outside option, these constraints are identical to the participation constraints and therefore redundant. Deleting those incentive constraints makes no difference to the set of feasible prices (and, more generally, sales mechanisms) for the non-students.

This logic generalizes to a broad class of mechanism design problems: namely, those in which the agents have access to an outside option. I show that whenever there is a set of types who receive the outside option in the benchmark mechanism, and who would do better than the outside option if that set was directly observable, and appropriate monotonicity and single-crossing conditions apply, there is an evidence structure under which no type is worse off and some types are better off. This evidence structure can also arise in equilibrium when evidence is chosen endogenously.

The result that an equilibrium with evidence may Pareto dominate the benchmark equilibrium with no evidence contrasts with results from the disclosure literature (Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986). The typical result in that literature is that informed parties with access to arbitrarily precise evidence are forced into full disclosure in any equilibrium. In the buyer-monopolist context this would lead to zero rents for all buyers. In my model, partial disclosures are possible in equilibrium, which make excluded buyers strictly better off without affecting the payoff received by the remaining buyers. The differences between my model and the models used in the disclosure literature are discussed in more detail below.

1.1 Related literature

The most closely related existing papers are Sher and Vohra (2015) and McAdams (2011), both of which consider a model of a monopolist facing a buyer with evidence. Sher and Vohra characterize optimal mechanisms given arbitrary sets of evidence that may be available to the buyer. They show that optimal mechanisms with evidence differ from the cheap-talk case in several ways. In particular, evidence may induce second degree price discrimination, even though buyers have unit valuations, and may lead to non-monotone allocations. Unlike in this paper, Sher and Vohra do not consider welfare implications or endogenous evidence acquisition. Moreover, my results generalize to a broad class of adverse selection environments whereas Sher and Vohra's results are specific to a buyer-seller relationship.

McAdams considers a model with endogenous evidence; however the model is different in several ways. Firstly, the timing is opposite: the monopolist commits to a mechanism before the buyer acquires evidence (which, in McAdams' model, is costly to acquire). Secondly, the buyer's choice of evidence is restricted to either exact evidence of the buyer's value or no evidence at all. In contrast, I allow for arbitrary partial evidence. McAdams shows that, in this model, aggregate welfare can be non-monotone in the cost of evidence. As with Sher and Vohra, McAdams' analysis is restricted to a buyer-seller model.

More broadly, this paper is related to the literature on mechanism design with evidence. Green and Laffont (1986), Bull and Watson (2004; 2007) and Koessler and Perez-Richet (2014) consider when social choice functions can be partially implemented with evidence, and to what extent analogues of the revelation principle apply. Green and Laffont, as well as Bull and Watson, find conditions under which direct mechanisms in which the agent discloses all available evidence are without loss of generality: these conditions apply in the model considered in this paper.

Koessler and Perez-Richet characterize social choice functions which are (partially) implementable with *evidence-based* mechanisms: mechanisms which implement an outcome consistent with the social choice function and presented evidence both on- and off-path. Kartik and Tercieux (2012) and Ben-Porath and Lipman (2012) consider full implementation, in the sense of Maskin (1999).

In my paper, many of the issues considered in the literature on partial implementation with evidence do not arise. In particular, I allow agents to costlessly disclose all evidence they obtain: this implies a condition known variously as Normality (Bull and Watson, 2007) or the Nested Range Condition (Green and Laffont, 1986) under which restricting attention to mechanisms in which all available evidence is disclosed — and, in addition, agents truthfully report their private information — is without loss of generality.

As noted above, my results contrast with the literature on disclosure (Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986). In the typical model found in the disclosure literature, a seller has private information about the quality of a good for sale. The seller can make truthful, possibly partial, disclosures about the quality of the good. Higher quality goods command a higher price. The main result is that full disclosure is the unique equilibrium: for any partial disclosure policy, if the true quality of the good is higher than the expectation given the partial disclosure, the seller does better by making a further disclosure. Any policy involving less than full disclosure thus 'unravels' as the highest types deviate.

It is clear from the intuition that this unravelling result would continue to hold if the seller could choose what evidence to acquire rather than simply what to disclose. Likewise, an equilibrium with partial pooling would continue to exist in my example if the buyer could make an overt disclosure *before* the seller commits to a mechanism (the buyer has no incentive to make any further disclosure, as he receives the lowest price consistent with the disclosure made in the equilibrium constructed below). Rather than the timing, the key difference is that in the disclosure models, the 'receiver' treats the 'sender' as if his type was the expected value (or, at least, some intermediate value) given the receiver's posterior belief after observing a disclosure. In my model, as the construction below will make clear, for some signal structures, the sender is treated as if his type were the lowest possible given the disclosure. Since all types of the sender, in both models, have monotonic preferences over the receiver's belief, this leads to unravelling in the disclosure models but not in my model.

A large literature, following the initial work on disclosure, shows that the unravelling result can be avoided by various plausible modifications to the basic model. There are two main strands. In the first strand disclosure is costly (representative papers include Jovanovic, 1982; Verrechia, 1983; 1990; Dye, 1986; Lanen and Verrechia, 1987). In the second strand there is uncertainty about whether the sender is able to make a disclosure (Dye, 1985; Jung and Kwon, 1988). In both strands, it is no longer possible to make a worst-case inference about a seller who does not fully disclose the quality. The mechanism allowing for partial disclosure in my model is unrelated to these arguments.

Giovannoni and Seidmann (2007) also provide conditions under which an equilibrium with partial pooling exists in a game with verifiable disclosure. In their model, actions are one-dimensional and the difference between the sender's and the receiver's ideal action, for a given state, can be characterized by a real valued bias function. There is an equilibrium with pooling on some subset of the state space if the direction of the bias switches on that subset and the receiver's ex-ante preferred action on that subset lies between the highest ex-post preferred action for states where the bias is negative and the lowest ex-post preferred action for states where the bias is positive.

In my example this property does not hold: if we restrict attention to posted prices, the sender (i.e. the buyer) always prefers lower prices and the receiver (i.e. the monopolist) prefers ex-post to charge the buyer's valuation, so that the 'bias' is always in the same direction. Moreover, the example is not a case of Giovannoni and Seidmann's framework since the receiver's action space is multi-dimensional. However, a common theme is that there is no feasible direction in which the sender would like the receiver to update: in Giovannoni and Seidmann's partial revelation result the sender is treated as an intermediate type but would like the receiver's belief to move in the opposite direction of the truth; in my construction the sender is treated as the most favorable type given the disclosed subset.

In a recent paper, Hagenbach, Koessler and Perez-Richet (2014) give sufficient conditions for existence of a full disclosure equilibrium – an equilibrium in which all evidence is disclosed – in a general setting. Their conditions are sufficient for a full disclosure equilibrium to exist, but do not ensure that it is unique. In fact, my results give an example where their conditions apply yet a partial disclosure equilibrium also exists.

This paper is also related to the voluminous literature on third degree price discrimination (see Armstrong, 2006, for a recent survey). Most closely related is the strand of the literature concerned with the welfare implications of third degree price discrimination. Schmalensee (1981), Varian (1985) and Schwartz (1990) establish, in increasingly general models, that a necessary condition for third degree price discrimination to improve welfare over a uniform monopoly price benchmark is that output increases. This necessary condition continues to hold with evidence. However, in contrast to these papers I focus on consumer surplus. It is quite possible for third degree price discrimination to increase total welfare while reducing consumer surplus: an extreme example is perfect price discrimination. Moreover, my model shows how evidence structures which increase consumer surplus can arise endogenously, while these papers consider exogenous observable information.

In a recent paper, Bergemann, Brooks and Morris (2015) characterize the set of pairs of consumer and producer surplus that can be achieved in equilibrium for any possible observable information the monopolist may have about the buyer's valuation. In comparison to Bergemann, Brooks and Morris my model covers a broader range of environments and my results in the specific case that they study may be more appropriate when evidence is endogenous.

Finally, the mechanics of the model are related to the recent literature on positive selection in dynamic mechanism design (Board and Pycia, 2014; Tirole, 2015). In this literature, a mechanism designer with limited commitment faces an agent in a dynamic interaction over time. Positive selection means that 'low types' – those with whom less aggregate surplus can be created – exit over time. This allows the designer to adhere to the full commitment mechanism, since positive selection makes deviations from the full commitment mechanism less attractive over time.

In my model, we can think of the set of types of the buyer that acquire no evidence as the 'general market'. By acquiring evidence, low types exit the general market and participate in a separate market which requires evidence to enter. Since the low types exit there is positive selection into the general market. The argument that the price in the general market is unaffected is analogous to the argument, in the positive selection literature, that the principal can commit to the optimal static mechanism.

2 Buyer-Monopolist Model

In this section I consider a model in which a buyer with a private unit valuation for a good (he) purchases from a monopolist (she). Prior to interacting with the monopolist, the buyer may have access to some evidence about his value. Given the available evidence, the monopolist commits to a revenue-maximizing mechanism. I show that there exists an evidence structure under which each type of the buyer is better off than if no evidence was available to him. Moreover, the equilibrium is an interim Pareto improvement: no type of the buyer is worse off while some types of the buyer are strictly better off.

The basic model is standard: the buyer has a private unit valuation, $v \in [0,1]$, for a good supplied by a monopolist. The buyer's value is drawn from a distribution, $F(v)$, with density $f(v)$. I make the following assumption:

Assumption 1 *The density, $f(v)$ is strictly positive on $[0,1]$ and the virtual value*

$$v - \frac{1 - F(v)}{f(v)}$$

is strictly increasing in v.

The buyer's ex-post payoff if he receives the good and pays a price p to the monopolist is $v - p$. The monopolist can supply the good at zero cost and is a revenue maximizer.

2.1 Exogenous Evidence

Before approaching the monopolist, the buyer may have access to evidence in the form of statements, "$v \in E$" for any Lebesgue measurable subset, E, of $[0,1]$ such that $v \in E$. The first requirement is an innocuous technical condition. The second requirement reflects the fact that evidence statements cannot be false (though they need not be the whole truth). Let $\mathcal{L}$ denote the Lebesgue measurable subsets of $[0,1]$

An *evidence structure* is a correspondence $\mathcal{E} : [0,1] \rightrightarrows \mathcal{L}$.

In addition, each type is always capable of sending no evidence; that is, the statement "$v \in [0,1]$". The evidence available to type v given evidence structure $\mathcal{E}$ is $\mathcal{E}(v) \cup [0,1]$.

After the buyer has acquired evidence, the seller commits to a mechanism. Since the buyer is in possesion of evidence it is not *a priori* without loss of generality to restrict attention to direct mechanisms. However, under the assumption that the buyer can always send *all* evidence statements that are available to him, the evidence structure satisfies a condition called *Normality* proposed by Bull and Watson (2007; and related to the *Nested Range Condition* in Green and Laffont (1986)).

Normality requires that each type, v, can be associated with a 'maximal' evidence statement, such that if some other type w can present type v's maximal evidence, then type w can present all evidence that type v can present. Normality is automatically satisfied if all types can present all the evidence available to them. That is, there are no costs or time constraints associated with presenting evidence.

Bull and Watson show that under Normality we can restrict attention to mechanisms in which the message space is the type space together with the set of possible evidence statements. All types truthfully report their type and present all evidence that is available to them.

Proposition 1 (Bull and Watson, 2007): *If the evidence structure satisfies Normality, then a social choice function is (partially) implementable if and only if it is implementable with truthful cheap talk messages and full evidence disclosure.*

For a given evidence structure an optimal mechanism for the seller solves:

$$max_{\alpha(v),p(v)} \int_0^1 p(v)dF(v)$$

$$s.t.\ \alpha(v)v - p(v) \geq \alpha(v')v - p(v')\ \forall (v,v') : \mathcal{E}(v') \subset \mathcal{E}(v)$$

and

$$\alpha(v)v - p(v) \geq 0\ \forall v.$$

We can construct an evidence structure under which no type of the buyer does worse, and some types do better, than they would if no evidence was available and which, moreover, is efficient. In the absence of evidence, the optimal mechanism for the monopolist is a posted price, $p_1^* \in (0,1)$. The equilibrium payoff is $v - p_1^*$ if $v \geq p_1^*$ and 0 otherwise. Let

$$u^*(v) = \begin{cases} v - p_1^* & v \geq p_1^*, \\ 0 & v < p_1^*. \end{cases}$$

Given an arbitrary mechanism, let $u(v)$ be the payoff to type v of the buyer.

Proposition 2 *For any distribution of types, $F(v)$, satisfying Assumption 1 there exists an evidence structure and an optimal mechanism given that evidence structure with:*

1. *$u(v) \geq u^*(v)$ for all $v \in [0,1]$.*
2. *$u(v) > u^*(v)$ for all $v \in [0,p_1^*]$ except for a set of measure zero.*

Moreoever, any optimal mechanism given that evidence structure induces an interim Pareto-improvement.

We prove proposition 2 by construction. Recall that p_1^* is the optimal monopoly price when no evidence is available. Recursively, define

$$p_k^* = min\{argmax_{v\in[0,1]} v(1 - F(v|v < p_{k-1}^*))\}.$$

Lemma 1 *p_k^* is well defined and the sequence $\{p_k^*\}_{k=1}^{\infty}$ converges to zero.*

The proof of Lemma 1 is contained in the appendix.

For $v \in (0, p_1^*)$ let

$$A(v) \equiv [p_k^*, p_{k-1}^*)$$

for k such that $v \in [p_k^*, p_{k-1}^*)$.

The evidence structure is as follows:

$$\mathcal{E}(v) = \begin{cases} \{0\} & v = 0, \\ [0,1] & v \geq p_1^*, \\ A(v) & 0 < v < p_1^*. \end{cases}$$

I refer to this evidence structure as $\mathcal{A}$.

Given this evidence structure, an optimal mechanism for the monopolist is to charge 0 to buyers presenting $\{0\}$, p_k^* to buyers presenting $[p_k^*, p_{k-1}^*)$ and p_1^* to buyers presenting no evidence. Without loss of generality, the monopolist does not sell to the buyer if presented with evidence statements outside $\cup_v \mathcal{E}(v)$.

Since the no evidence price is the highest, all available evidence is presented in equilibrium. To see that the mechanism is optimal, note that p_1^* continues to be the optimal price when the seller faces the distribution $F(v|v \geq p_1^*)$. Clearly a lower price cannot be optimal. If a higher price generated higher revenue, it would also generate higher revenue against the distribution, $F(v)$. This follows because with both p_1^* and any price $p' > p_1^*$ no buyers with values below p_1^* purchase. These buyers are therefore irrelevant to a revenue comparison between p_1^* and $p' > p_1^*$. By the same argument the optimal mechanism is still a posted price (that is, $\alpha(v) \in \{0, 1\}$).

Similarly, the optimal price given $F(v|[(p_k^*, p_{k-1}^*))$ is the same as the optimal price given $F(v|v < p_{k-1}^*)$.

Since no type pays more than p_1^* and all types purchase, $u(v) \geq u^*(v)$ for all $v \in [0,1]$. For any $v < p_1^*$, $u^*(v) = 0$, whereas $u^*(v) > 0$ for almost all $v \in (0, p_1^*)$ in the constructed equilibrium.

Given Assumption 1, p_1^* is the unique optimal price on $[p_1^*.1]$ (or any subset of $[0,1]$ including $[p_1^*, 1]$) so that, given $\mathcal{A}$, p_1^* is the unique optimal no-evidence price. Moreover, any optimal price on $[p_2^*, p_1^*)$ must be strictly lower than p_1^*. It follows that *any* optimal mechanism given evidence strucutre $\mathcal{A}$ induces an interim Pareto-improvement.

The observation that truncating the distribution below p_1^* does not affect the optimal price is key to the results in this paper: if buyers presenting evidence that their value is below p_1^* receive a lower price, this does not imply that buyers with values above p_1^* receive a higher price, allowing for an interim Pareto improvement. The key property is that buyers below p_1^* take their outside option in the equilibrium without evidence. Because of this, the incentive constraints between high value buyers and low value buyers are redundant: they are satisfied whenever the high value buyer's participation constraints are satisfied.

Note, however that this feature is not sufficient for the existence of an evidence structure that induces an interim Pareto-improvement (see Example 5). Two properties of this environment ensure that this feature is sufficient in this case. In the first place, the principal's preferences over outcomes (here, pairs of transfers and allocation probabilities) do not depend on the agent's type. Secondly, a single-crossing condition on the agent's preferences is satisfied. Either of these properties is sufficient to ensure that when evidence is available, the principal does not want to change the mechanism offered to the previously included set of types solely because incentive constraints corresponding to imitating types in the previously excluded set have been removed. In the final section, I generalize this logic to show that evidence can lead to an interim Pareto improvement in a broad class of environments.

2.2 Evidence Provided by a Planner

In many environments, evidence is regulated or made available by a central authority with the interests of the agents who acquire evidence in mind. For example: student ID cards are provided by a university, professional organizations provide certifications to their members, governments provide senior citizen cards. In this section, I consider a variant of the model used in the previous section with the difference that a planner — concerned with maximizing consumer surplus — selects which evidence statements the buyer can choose to obtain, or equivalently which statements can be contracted upon. I show that there exists a policy for the planner such that the interim payoffs from the equilibrium described above are achieved (almost surely) in any iteratively admissible equilibrium. Moreover, under the same policy, the buyer is better off compared to the no-evidence benchmark in any equilibrium, whether admissible or not.

The planner's decision is observable to both the buyer and the monopolist. Formally, the planner chooses a measureable correspondence

$$\chi : v \rightrightarrows 2^{\mathcal{L}}$$

Sastifying $v \in E$ for all $E \in \chi(v)$. The buyer then chooses an evidence acquisition strategy, $\mathcal{E}$, (a set of evidence statements for each type) with the restriction $\mathcal{E}(v) \subset \chi(v)$. The game then continues as in the previous section.

To summarize, the timing is:

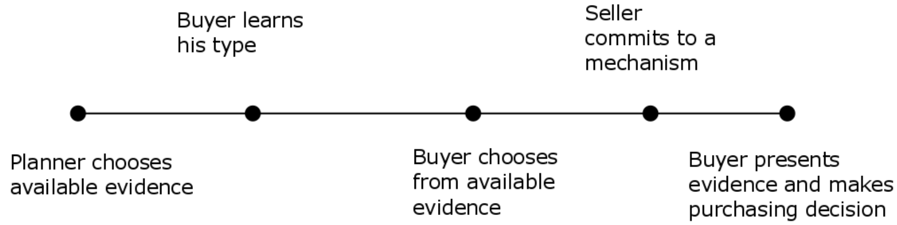

The main result of this section shows how a planner can ensure an interim Pareto-improvement across all equilibria by restricting the available evidence, under a weak refinement: admissibility. Admissibility simply requires that weakly dominated strategies are not played.

It will be clear that, for this result, it is only necessary that the planner choose a partition of the state space and allow the buyer to acquire evidence proving which element of the partition his type is in, and disallow the buyer to acquire any other evidence. Note that the planner does not choose the evidence-acquisition strategy, but simply restricts the available evidence, from which the buyer is then free to choose.

Proposition 3 *There is a policy for the planner, and a continuation equilibrium, such that:*

1. $u(v) \geq u^*(v)$ *for all* $v \in [0, 1]$
2. $u(v) > v^*(v)$ *for almost all* $0 < v < p_1^*$.

Moreover, given the same policy, every admissible continuation equilibrium induces an interim Pareto-improvement over the no-evidence benchmark.

Proof: Let

$$\chi(v) = \begin{cases} [0,1] & v \in [p_1^*, 1], \\ \{0\} & v = 0, \\ A(v) & \text{otherwise}, \end{cases}$$

where $A(v)$ is defined as in the previous section. There is an equilibrium satisfying the proposition, since if all types of the buyer acquire evidence, an optimal mechanism for the monopolist is to charge p_1^* to buyers presenting no evidence, p_k^* to buyers presenting A_k and 0 to buyers presenting $\{0\}$. Given this, it is clearly optimal for each type of the buyer to acquire the available evidence.

To see that no other equilibrium evidence structure satisfies admissibility, we note that not acquiring all available evidence is weakly dominated. Indeed, the buyer cannot be made worse off by acquiring $\chi(v)$ since he can always choose not to present it. On the other hand, given the monopolist's equilibrium mechanism outlined above it is a strict best response for (almost every type of) the buyer to acquire all available evidence in $\chi(v)$. It follows that (almost) no other equilibrium evidence structure satisfies admissibility.

To see that the constructed equilibrium satisfies admissibility note that acquiring full evidence can never be (weakly) dominated, since the buyer is always free not to present the evidence he has acquired. Given that the unique admissible equilibrium evidence structure is $\mathcal{A}$, the proposition is satisfied. The argument is identical to the one made for Proposition 2.

The same argument applies in the general case: whenever an equilibrium which Pareto-dominates the no-evidence benchmark exists, the planner can ensure that that equilibrium is unique under admissibility by suitably restricting the available evidence.

If the buyer is allowed to use a random evidence-acquisition strategy, the planner can do even better. Bergemann, Brooks and Morris (2015) consider the following question: what is the set of pairs of consumer and producer surplus that are achieved in equilibrium for any arbitrary information a monopolist can observe about the buyer's value? They construct an information structure that maximizes consumer surplus given a best response by the monopolist (henceforth referred to as the BBM information structure).

Under the BBM information structure, after any signal, the monopolist is indifferent between charging the value of the lowest type who generates that signal with positive probability and charging the no-information monopoly price. All types of the buyer are served, hence the equilibrium is efficient. On the other hand, since the no-information monopoly price is also optimal on each segment, the monopolist's profit is the same as with no information. Since this is a lower bound on the monopolist's profit under *any* information structure, consumer surplus is maximized.

Under the (deterministic) evidence structure that I construct, the monopolist makes a strictly higher profit than under the no-evidence benchmark; hence consumer surplus is lower than under the BBM information structure. The BBM information structure is consistent with an evidence acquisition strategy in which every type acquires exactly one evidence statement (though each type must randomize over several available statements). Since the buyer always has evidence the monopolist optimally excludes the buyer if he presents no evidence. It follows that the acquired evidence is always presented given any individually rational price so that evidence functions identically to observable information in this context.

Although the BBM information structure can lead to maximal consumer surplus, providing the BBM information structure is unappealing for a planner for two reasons. Firstly, it requires the buyer to play a specific mixed strategy, and in fact one that is weakly dominated, while the deterministic equilibrium that I construct only requires the buyer to play the unique admissible strategy. Secondly, given that the buyer's information acquisition strategy is consistent with the BBM information structure, there is another optimal mechanism for the monopolist that leaves the buyer with low surplus (under the BBM information structure the monopolist is indifferent between charging the lowest value consistent with each signal and charging the no-information monopoly price after each signal).

2.3 Endogenous evidence and disclosure

In this section I show that the key result (the possibility of an interim Pareto improvement) is robust to some features of the timing of the game. In the benchmark case, we assumed that a given evidence strucutre was exogenously available to the agent. When discussing policy implementation, we noted that, if a planner restricts the feasible set of evidence statements and then allows the agent to acquire, or not acquire, evidence as he chooses then the same evidence structure arises in equilibrium and is essentially unique under admissibility.

Even with no planner, however, the same equilibrium continues to exist in an environment where the buyer can choose what evidence he acquires (from an unrestricted set of feasible statements). Consider a model with the following timing:

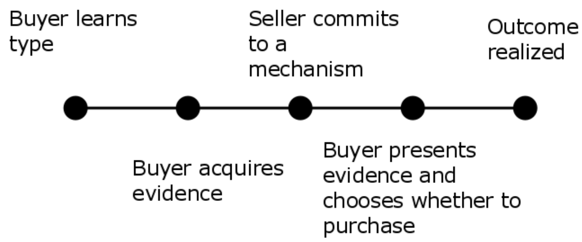

The buyer observes the seller's choice of mechanism before participating in the mechanism. The seller does not observe the buyer's evidence acquisition strategy before committing to a mechanism.

The solution concept is based on Perfect Bayesian Equilibrium (Fudenberg and Tirole, 1991). An equilibrium consists of a mechanism chosen by the seller, a reporting strategy within the mechanism for the buyer and an evidence-acquisition strategy for the buyer such that:

1. The reporting strategy is optimal for each type of the buyer given the mechanism.

2. Any mechanism chosen with positive probability is optimal given correct beliefs about the evidence acquisition strategy and the prior on the buyer's type.

3. The buyer's evidence acquisition strategy is optimal given correct beliefs about the seller's strategy and a plan to act optimally in the mechanism.

The same outcome continues to obtain in an equilibrium as follows: given that the seller chooses the same mechanism reporting truthfully and presenting all evidence is optimal. The mechanism is optimal given that the same evidence structure is chosen. Without loss of generality the seller offers the outside option when presented with any unexpected (off-path) evidence (since in equilibrium she correctly puts zero probability on the event that such evidence is available). Given this, acquiring evidence structure $\mathcal{A}$ is optimal for the buyer.

It will be clear that this equilibrium coexists with many other equilibrium. In fact, *any* equilibrium that consists of an optimal mechanism given some evidence structure for the seller, and that choice of evidence structure for the buyer is an equilibrium. This is because the seller without loss of generality ignores any off path evidence, so that there is no incentive to deviate to acquiring more evidence than is used in any equilibrium. Conversely, it is weakly dominated not to acquire all evidence used on path (as suggested in section 2.2) so that there is no incentive to acquire less evidence than is used in the given equilibrium. It will also be clear for the same argument that the only admissible equilibrium is the one in which all possible evidence is acquired and the seller engages in perfect price discrimination.

This highlights the role of policy: with no restrictions on the available evidence and where arbitrary evidence can be acquired, the unique admissible equilibrium leads to first-degree price discrimination: a very bad outcome from the point of view of the buyer. By restricting the evidence that can be acquired or contracted upon the planner can ensure a better outcome for the agent (and also, as we saw, an efficient outcome). This does not require the planner to actively select or provide evidence for the buyer: the logic of admissibility ensures that with suitable restrictions the desired evidence structure will be endogenously chosen in equilibrium.

Another alternative timing illuminates the connection, and contrast, between my model and the canonical disclosure models (e.g. Grossman, 1981; Milgrom, 1981). In this timing we assume that the buyer acquires evidence and then presents it to the seller *before* the seller commits to a mechanism. The timing is:

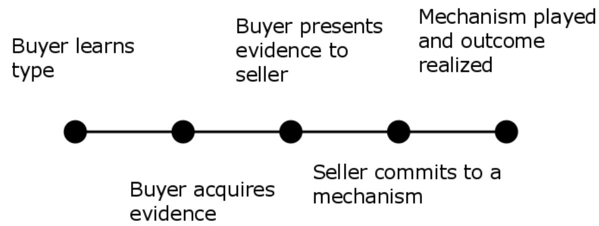

The seller observes any evidence that the buyer has chosen to present (but not the evidence-acquisition strategy) at the time she commits to the mechanism. The buyer observes the choice of mechanism before participating in it.

Again, the equilibrium concept is based on Perfect Bayesian Equilibrium. An equilibrium consists of, for each type of the buyer, a (history dependent) choice of what evidence to acquire, what evidence to present and how to report in the mechanism, and for the seller a (history dependent choice of mechanism) such that:

1. All choices, both on- and off-path are optimal given beliefs
2. Beliefs are correct on-path
3. The seller's beliefs about the buyers type, when choosing a mechanism, are consistent with any evidence presented. That is, if the buyer presents evidence statement "$v \in E$", then the seller's beliefs put zero probability on types of the buyer outside E.

With this timing the game shares some features with both disclosure models and models of interim Bayesian persuasion (Perez-Richet, 2014). The agent would like to persuade the seller that his value for the good is low, and persuasion takes place once the buyer already knows his value, unlike in the classic Bayesian persuasion literature (Kamenica and Gentzkow, 2011). Nevertheless, there exist equilibria that do not result in full disclosure, including an equilibrium leading to the same outcome as evidence structure $\mathcal{A}$.

We will see that an equilibrium exists in which the buyer chooses to acquire evidence structure $\mathcal{A}$, the buyer chooses to present all acquired evidence to the seller, the seller makes a take it or leave it offer of p_k^* for the relevant k, and the buyer accepts the offer. That it is a best response to accept the offer is straightforward, and that the offer p_k^* is a best response to the on-path belief that $v \in [p_k^*, p_{k-1}^*)$ follows from the discussion earlier in the paper.

To show that the evidence-acquisition and presentation strategies form part of an equilibrium we must consider the seller's beliefs when presented with off-path evidence. Consider the following off-path beliefs for out of equilibrium evidence statements (or sets of statements) E' (if E' is in fact a set of subsets, consider the following to apply to their intersection — which must include the deviating type and hence is nonempty). Let $\mu'(E')$ denote the seller's off-path belief on observing evidence statement E' and $\mu(\cdot)$ denote the measure on μ induced by the CDF F:

$$\mu'(E') = \begin{cases} sup_{k:p_k^* \in E'} p_k^* & \exists k : p_k^* \in E' \\ \mu(v|E') & otherwise \end{cases}$$

In the first case the seller's best response is to charge $\mu'(E') \geq p_k^*$ for the relevant k so that the buyer would be worse off making this deviation. In the second case, the seller's best response is to charge a price at least as high as $inf E' > p_k^*$, so that again the deviating type of the buyer is worse off.

Again, this equilibrium coexists with many others. For example, full disclosure is also an equilibrium. However, the existence of an equilibrium with partial disclosure seems on the surface to contradict the typical results of the disclosure literature, under which when complete and costless hard evidence is available, and preferences are monotone in a receiver's action (here think of price) the unique equilibrium is full disclosure.

In Milgrom's (1981) disclosure model, for example, disclosures can be made by a seller about the quality of a good that she has for sale. A buyer demands a certain quantity of the good, depending on the expected quality, and every type of the seller prefers to sell a higher quality. The unique sequential equilibrium is full disclosure because, whenever partial disclosures are made, unless the outcomes is the same as what would occur under full disclosure, some type of the seller sells a lower quantity than would occur is she fully disclosed her type. This follows because the quantity sold depends on the buyer's expectation of the quantity and the expected expectation is preserved by any disclosure policy (by the law of iterated expectations).

In contrast, in my model, different evidence disclosure strategies can lead to the same expectation over expected type, but to a different expectation over prices. For example, disclosing only "$v \in [p_1^*, 1]$" leads to a price of p_1^*, but fully disclosing the type on this set leads to an expected price of $E[v|[p_1^*, 1]] > p_1^*$. It is therefore possible that, for some evidence structures, any further disclosures must harm *all* types. $\mathcal{A}$ is an example of such a structure.

An underlying difference is that in my model the sender (here, the buyer) has evidence about his own payoffs while in Milgrom's model the sender (there, the seller) has evidence about the receiver's payoffs. This means that in my model, once enough evidence is available to overcome inefficiency in the revenue-maximizing mechanism, disclosing more evidence simply reduces the sender's information rents. This same effect does not hold in Milgrom's model (where revealing more evidence has the effect of 'redistributing' surplus across sender types).

3 Generalization

In the first section we saw that the availability of evidence can support an equilibrium that interim Pareto-dominates the benchmark where no evidence is available, in a model of a buyer with an unknown unit valuation purchasing from a monopolist. In this section I show that the result is far more general. In a general screening model where the agent has access to an outside option, and both the agent's and principal's preferences are monotone in allocations and transfers I show that there is an equilibrium with evidence that interim Pareto-dominates the benchmark equilibrium with no evidence whenever:

1. Some types of the agent are excluded (take their outside option) in the benchmark mechanism.
2. If the type space was restricted to the excluded types, the payoff to some type in an optimal mechanism is greater than the payoff from the outside option.
3. The agent's preferences satisfy a single-crossing condition.

The third condition can also be replaced by the condition that the principal's preferences over allocations do not depend on the agent's type. These conditions hold in a broad range of economically relevant environments. For example, almost any model in which the principal is a profit-maximizing monopolist whose costs of production do not depend on the buyer's private information satisfies the conditions. The conditions also apply to the classic monopolistic insurance model of Stiglitz (1977), which inherently incorporates type-dependence in the principal's payoffs.

Before stating the result formally, we introduce some notation:

A principal (she) offers a mechanism to an agent (he). The agent's type $t \in T$ is private information and is distributed according to a probability measure $\pi \in \Delta(T)$. T is assumed to be a complete, compact, metric space: any statements about continuity are with respect to the given metric. The set of feasible alternatives are of the form $(a, p) \in A \times P \subset \mathbb{R}^+ \times \mathbb{R}$, where a is interpreted as an allocation and p as a transfer.

The principal and agent have state-dependent, continuous preferences $u : A \times P \times T \to \mathbb{R}$ and $v : A \times P \times T \to \mathbb{R}$, respectively. Both are expected utility maximizers.

We assume that $u((a, p), t)$ is weakly increasing in a and strictly decreasing in p, while $v((a, p), t)$ is weakly decreasing in a and strictly increasing in p.

The principal commits to a mechanism. A mechanism is a message space, M, together with an allocation function $g : M \to \Delta(A \times P)$. In the benchmark case where no evidence is available, restricting attention to direct mechanisms is without loss of generality, so that $M = T$. With slight abuse of notation, we will sometimes refer to an outcome function as a mechanism. Given that the message space is fixed, this is without loss of generality. A mechanism must be incentive compatible, that is:

$$u(g(t), t) \geq u(g(t'), t) \;\; \forall t, t' \in T.$$

Throughout this section we maintain the assumption that the agent has access to the outside option $(0, 0)$. The outside option is always available to the agent. The mechanism offered by the principal must respect the participation constraint:

$$u(g(t), t) \geq u((0, 0), t) \;\; \forall t \in T.$$

When the agent has acquired evidence there is an exogenous meaning to certain messages. In particular, let $E(t)$ be the set of evidence statements available to type t. If $E(t) \subsetneq E(t')$, then the principal does not have to respect incentive constraints from t to t'. This follows given Normality (discussed above) which is automatically satisfied given that there is no cost to presenting all acquired evidence. Normality implies that all acquired evidence is presented, without loss of generality. Then if $E(t) \subsetneq E(t')$, given that t' discloses all evidence in $E(t')$ the mechanism designer knows that the agent's type is not $E(t')$ when his true type is t. Conversely, given that incentive constraints between types who can imitate each other are satisfied, the agent has no incentive not to present all available evidence.[2] The set of incentive constraints becomes

$$u(g(t),t) \geq u(g(t'),t) \;\; \forall t,t' : E(t) \supset E(t').$$

For a subset of types $\tilde{T} \subset T$, we define a restricted mechanism as a message space $M_{\tilde{T}} = \tilde{T}$ and an outcome function $g_{\tilde{T}} : \tilde{T} \to \Delta(A \times P)$.

Let $f : T \to \Delta(A \times P)$ satisfy

$$f \in argmax_g \int_T v(g(t),t)d\pi(t)$$

such that

$$u(g(t),t) \geq u(g(t'),t) \;\; \forall t,t' \in T,$$

$$u(g(t),t) \geq 0 \;\; \forall t \in T.$$

For $\tilde{T} \subset T$, let $f_{\tilde{T}} : \tilde{T} \to \Delta(A \times P)$ satisfy

$$f_{\tilde{T}} \in argmax_g \int_{\tilde{T}} v(g_{\tilde{T}}(t),t)d\pi(t)$$

such that

$$u(g_{\tilde{T}}(t),t) \geq u(g_{\tilde{T}}(t'),t) \;\; \forall t,t' \in \tilde{T},$$

$$u(g_{\tilde{T}}(t),t) \geq u((0,0),t) \;\; \forall t \in \tilde{T}.$$

That is, $f_{\tilde{T}}$ is a restricted mechanism that would be offered by the principal if $t \in \tilde{T}$ were directly observable. Recall that, unlike evidence, the principal sees directly observable information whether or not the agent chooses to reveal it.

With slight abuse of notation, I will also use (a,p) to denote a mixture, $(a,p) \in \Delta(A \times P)$ and $u((a,p),t)$ to denote the expected utility. For a mechanism, g, let $a_g(t)$ denote the marginal distribution of $g(t)$ on A. For $a', a \in \Delta(A)$ we say write $a' \succ a$ if a' first order stochastically dominates a.

We can now state the general result:

Proposition 4 *Suppose that there exists a set of types $\tilde{T} \subset T$ such that*

1. $f(t) = (0,0) \Leftrightarrow \tau \in \tilde{T}$,
2. $u(f_{\tilde{T}}(t),t) > u((0,0),t)$ *for some* $t \in \tilde{T}$, *and, moreover:*

[2] To see why this argument fails when Normality is not satisfied see Bull and Watson (2007) and Green and Laffont (1986)

3. *Either:*

   (a) *Whenever* $a = marg_A(a,p)$, $a' = marg_A(a',p')$ *and* $a' \succ a$ *and* $t' > t$

$$u((a',p'),t) \geq u((a,p),t) \Rightarrow u((a',p'),t') > u((a,p),t')$$

   Or:

   (b) $v : A \times T \to \mathbb{R}$ *does not depend on* t.

Then there exists an equilibrium with evidence that interim Pareto-dominates the benchmark mechanism f.

Proof: Suppose that all types in $\tilde{T}$ have evidence $E(t) = \{\tilde{T}\}$ and no other types have any evidence. We first show that, given this evidence structure, there is an optimal mechanism with evidence in which no types in $T \setminus \tilde{T}$ are made worse off (given condition 1 and the participation constraints, no type in $\tilde{T}$ can be worse off). We show this by contradiction. Suppose that for all optimal mechanisms with evidence, $\hat{g}$, some type in $T \setminus \tilde{T}$ is worse off compared to f. Then it must be that

$$\int_{T\setminus\tilde{T}} v(\hat{g}(t),t)d\pi(t) > \int_{T\setminus\tilde{T}} v(f(t),t)d\pi(t).$$

Since, otherwise, the SCF:

$$\hat{g}'(t) = \begin{cases} f(t) & t \in T \setminus \tilde{T} \\ \hat{g}(t) & t \in \tilde{T} \end{cases}$$

is incentive compatible with evidence and preferred by the principal to $\hat{g}$, contradicting the hypothesis that $\hat{g}$ is optimal.

Given this, either $\hat{g}(t) = f(t)$ for all $t \in T \setminus \tilde{T}$ in some optimal mechanism with evidence or

$$\int_{T\setminus\tilde{T}} v(\hat{g}(t),t)d\pi(t) > \int_{T\setminus\tilde{T}} v(f(t),t)d\pi(t)$$

in every optimal mechanism with evidence, $\hat{g}$.

Suppose the latter. We will see that it leads to a contradiction. Let $\hat{g}''$ be defined by:

$$\hat{g}''(t) = \begin{cases} (0,0) & t \in \tilde{T} \\ \hat{g}(t) & t \in T \setminus \tilde{T} \end{cases}$$

If $\hat{g}$ is incentive compatible without evidence, this would contradict the hypothesis that f is an optimal mechanism without evidence. So suppose that it is not. The only incentive constraints that might be violated are from some $t \in \tilde{T}$ to some $t \in T \setminus \tilde{T}$, since the constraints from $T \setminus \tilde{T}$ to $\tilde{T}$ are redundant given the participation constraints.

We now proceed in two cases. For the first case assume that (3a) holds. We can show, as a consequence of this single-crossing property that the following holds:

Lemma 2 *If* $\hat{g}$ *is optimal, then*

$$u((0,0),t) \geq u(\hat{g}(t'),t)$$

for all $t \in \tilde{T}, t' \in T \setminus \tilde{T}$

Proof: See appendix.

But this contradicts the hypothesis that $\hat{g}''$ is not incentive compatible without evidence.

For the second case, suppose that (3b) holds. Let

$$\tilde{g}(t) = \begin{cases} \hat{g}(t) & t \in T \setminus \tilde{T} \\ argmax_{a \in \{\hat{g}(t): t \in T \setminus \tilde{T}\} \cup \{(0,0)\}} u(a,t) & t \in \tilde{T} \end{cases}$$

That is: we allow types in $\tilde{T}$ to choose either any alternative in the range of $\hat{g}$ over $T \setminus \tilde{T}$ or the outside option. By construction this mechanism is incentive compatible without evidence. It is also weakly preferred by the principal to $\hat{g}''$ as long as $v(\hat{g}(t)) \geq v((0,0))$ for all $t \in T \setminus \tilde{T}$ (note that under (3b) the principal's preferences are type independent, to emphasize this we drop the t argument in $v(\cdot)$).

In fact, this is without loss of optimality. Suppose that $v(\hat{g}(t)) < v((0,0))$ for some $t \in T \setminus \tilde{T}$. Let $T' \subset T \setminus \tilde{T}$ be the set of types for which this is true. T' is a strict subset of $T \setminus \tilde{T}$ or else the mechanism assigning $(0,0)$ to all types is optimal.

Let $\hat{g}'''$ be defined by:

$$\hat{g}'''(t) = \begin{cases} argmax_{a \in \{\hat{g}(t): t \in T \setminus (\tilde{T} \cup T')\} \cup \{(0,0)\}} u(a,t) & t \in T' \\ \hat{g}(t) & t \in T \setminus T' \end{cases}$$

Then $\hat{g}'''$ is incentive compatible with evidence and preferred by the principal to $\hat{g}$, contradicting the hypothesis that $\hat{g}$ is optimal with evidence.

By contradiction, in either case, we have shown that no type is made worse off (no type in $\tilde{T}$ can be worse off since they are held to their reservation payoff in f). We next show that some types are made better off. Suppose not. We have established that in some optimal mechanism with evidence $\hat{g}(t) = f(t)$ for all $t \in T \setminus \tilde{T}$. If no types in $\tilde{T}$ are better off then

$$u(\hat{g}(t), t) = u((0,0), t) \ \ \forall t \in \tilde{T}$$

By condition 2

$$\int_{\tilde{T}} v(f_{\tilde{T}}(t), t) d\pi(t) \geq \int_{\tilde{T}} (\hat{g}(t), t) d\pi(t)$$

and

$$u(f_{\tilde{T}}(t), t) > u((0,0), t)$$

for some $t \in \tilde{T}$. Let

$$\hat{g}'''(t) = \begin{cases} f_{\tilde{t}}(t) & t \in \tilde{T} \\ \hat{g}(t) & t \in T \setminus \tilde{T} \end{cases}$$

Then $\hat{g}'''$ is incentive compatible with evidence, since if $t \in \tilde{T}, t' \in T \setminus \tilde{T}$

$$u(\hat{g}'''(t), t) = u(f_{\tilde{T}}(t), t) \geq u((0,0), t) \geq u(\hat{g}(t'), t) = u(\hat{g}'''(t'), t)$$

Where the first inequality follows by IR of $f_{\tilde{T}}$ and the second inequality follows by IC of $\hat{g}$. Incentive constraints within $T \setminus \tilde{T}$ and within $\tilde{T}$ are satisfied by construction while incentive constraints from $T \setminus \tilde{T}$ to $\tilde{T}$ can be ignored given that all types $t \in \tilde{T}$ present evidence.

Since $\hat{g}'''$ is IC with evidence and is preferred to $\hat{g}$ by the principal, the optimality of $\hat{g}$ is contradicted. We conclude that $u(\hat{g}(t), t) > u((0,0), t)$ for some $t \in \tilde{T}$ in an optimal mechanism with evidence, as required.

3.1 Examples

The following examples demonstrate the broad applicability of conditions 1-3:

3.1.1 Second-degree price discrimination

In the leading example the monopolist faced constant (in fact zero) costs and the consumer valued at most one unit of the good. However, the general result applies to a much broader range of buyer-monopolist models. For concreteness, consider the following second-degree price discrimination model taken from Mussa and Rosen (1978):

Example 1 *The type space is $T = [\underline{\theta}, \overline{\theta}]$. The outcome space is the set of pairs $(q, p) \in \mathbb{R}^+ \times \mathbb{R}$ where $q \geq 0$ denotes a quality level and p a transfer from the buyer to the monopolist. The outside option is $(0, 0)$. The buyer's payoff is*

$$u((q,p), \theta) = \theta q - p$$

The monopolist's payoff is

$$v((q,p), \theta) = p - aq - bq^2$$

for parameters a and b. Types are uniformly distributed.

Mussa and Rosen show that the optimal mechanism without evidence involves $(q,p)(\theta) = (0,0)$ for $\theta \leq \theta^* \equiv (\overline{\theta} + a)/2$. Moreover, above θ^* the mechanism is fully separating and leaves positive payoffs to almost every type.

Since the mechanism design problem with $T = [\underline{\theta}, \theta^*]$ is simply a rescaled version of the original problem, is it clear that the optimal mechanism when $T = [\underline{\theta}, \theta^*]$ leaves positive payoffs to some type as long as a is not too large. So far we have seen that conditions 1 and 2 are satisfied.

Let $q' > q$ and $\theta' > \theta$. If $u((q', p'), \theta) \geq u((q, p), \theta)$ then

$$\theta(q' - q) \geq p' - p$$

$$\Rightarrow \theta'(q' - q) > p' - p$$

$$\Rightarrow u((q', p'), \theta') > u((q, p), \theta)$$

Since the agent is risk neutral, the result follows through when q', q are nondegenerate distributions and $q' \succ q$. Condition 3a is satisfied (it is also straightforward to see that condition 3b is satisfied, given that costs of production are type-independent).

Our result therefore implies that there is an equilibrium with evidence that delivers an interim Pareto-improvement in this environment.

3.1.2 Monopolistic insurance

A rather different environment in which the general result applies is an insurance market. The following example adapts the monopolistic insurance model of Stiglitz (1977):

Example 2 *The type space is $T = [\underline{t}, \bar{t}] \subset (0,1)$ where t is the probability of a loss. The outcome space is the set of pairs $(a,p) \in \mathbb{R}^+ \times \mathbb{R}$ where a is the gross amount paid by the insurer to the consumer if there is a loss, and p is a premium paid to the insurer in both states. The consumer's payoff is $u((a,p),t) = U(W-d-p+a)t + U(W-p)(1-t)$ where U is CARA. The insurer's payoff is $v((a,p),t) = p - a \cdot t$. t is distributed according to the CDF $F(t)$. The outside option is $(0,0)$. Assume the environment is such that a deterministic mechanism is optimal.*

Stiglitz shows that, under some parameterizations, some types take the outside option – always those types with the lowest probabilities of a loss (in the optimal deterministic mechanism). Typically, if the type space was restricted to those types some types would purchase insurance at a premium less than their willingness to pay. Conditions 1 and 2 are therefore satisfied.

To see that condition 3a is satisfied, let $a' \succ a$, $t' > t$ and

$$u((a', p'), t) \geq u((a, p), t) \quad (1)$$

Suppose that $E_{(a',p')}U(w-p') \geq E_{(a,p)}U(W-p)$. Then by CARA

$$E_{(a',p')}U(w-p'-d) \geq E_{(a,p)}U(W-p-d)$$

which implies

$$u((a', p'), t'') = t'' E_{(a',p')}U(W-p'-d+a') + (1-t'')E_{(a',p'))}U(W-p')$$

$$> t'' E_{(a,p)}U(W-p-d+a) + (1-t'')E_{(a,p)}U(W-p) = u((a,p), t'')$$

for any $t'' \in T$, where the strict inequality follows by $a' \succ a$. Conversely, suppose that $E_{p'}U(W-p') < E_p U(W-p)$. Then, by (1):

$$t(E_{(a',p')}U(W-p'-d+a') - E_{(a,p)}U(W-p-d+a))$$

$$\geq (1-t)(E_{(a,p)}U(W-p) - E_{(a',p')}U(W-p')) > 0$$

$$\Rightarrow t(E_{(a',p')}U(W-p'-d+a') - E_{(a,p)}U(W-p-d+a))$$
$$> (1-t)(E_{(a,p)}U(W-p) - E_{(a',p')}U(W-p'))$$
$$\Rightarrow u((a',p'),t') > u((a,p),t)$$

so that, in either case, condition 3a is satisfied.

Note that although we do need some assumption on the agent's vN-M utility to ensure that condition 3a is satisfied, we do not necessarily need to assume CARA preferences, this is simply done here to demonstrate that the condition is satisfied for *some* class of preferences.

Since the conditions of proposition 4 are satisfied, there is an equilibrium with evidence that interim Pareto-improves on the equilibrium without evidence is this monopolistic insurance market.

3.1.3 Tightness of Proposition 5

The remaining examples demonstrate that the conclusion of proposition 4 can fail if any of the conditions are relaxed.

Example 3 $T = [\frac{1}{2}, 1]$. $A = [0,1]$. *The principal's payoff is* $v((a,p),t) = p$, *the agent's payoff is* $u((a,p),t) = a \cdot t - p$. *Types are distributed uniformly.*

This is a version of the leading example with the type space truncated at the monopoly price. The optimal monopoly price with no evidence is $p = \frac{1}{2}$. The outcome is efficient and clearly no type can be better off as the principal will never charge less than $\frac{1}{2}$. In this example condition 1 fails: no type is excluded without evidence.

Example 4 $T = \{1,2\}$, $A = \mathbb{R}^+$. *The agent's payoff is* $u((a,p),t) = t \cdot a - p$, *the principal's payoff is* $p - \frac{1}{2}a^2$. $\pi(2) > \frac{1}{2}$

By standard arguments, the optimum has $(a,p)(2) = (2,4)$ and $(a,p)(1) = (0,0)$ (type 1 is not excluded if $\pi(2) < 2$). However, condition 2 is not satisfied, since if $t = 1$ is observable the optimal contract is $(1,1)$ and $u((1,1),1) = u((0,0),1) = 0$. We will show that for any evidence structure, type 1 does not recieve a positive payoff and type 2 does no better than in the mechanism without evidence.

There are three possible evidence structures that differ from no evidence: both types have evidence of their type, type 1 has evidence and type 2 does not, or type 2 has evidence and type 1 does not. In the first and third cases the mechanism designer can implement the efficient allocation and fully extract all surplus. In the second case, it is easy to verify that the no-evidence contract remains optimal. It follows, as claimed, that there is no equilibrium with evidence that interim Pareto-dominates the mechanism offered with no evidence.

Example 5 *The type space is* $T = \{t_1, t_2, t_3, t_4\}$, *the set of pure outcomes is* $\{a_0, a_1, a_2\}$: *for example* a_0 *is the outside option and* a_1, a_2 *are two different goods of which the agent demands at most one unit in total. Preferences are given by:*

$$u(a,t) - p$$

and

$$p - v(a,t)$$

respectively, where $u(\cdot)$ *and* $v(\cdot)$ *are given by:*

$u(\cdot)$	t_1	t_2	t_3	t_4
a_0	*0*	*0*	*0*	*0*
a_1	*1/2*	*1*	*3*	*3*
a_2	*6*	*6*	*5*	*5*

$v(\cdot)$	t_1	t_2	t_3	t_4
a_0	*0*	*0*	*0*	*0*
a_1	*0*	*0*	*0*	*0*
a_2	*-20*	*-20*	*0*	*0*

In this example, absent types t_1 and t_2, the principal would like to allocate a_2 for sure to t_3 and t_4 and charge $p_3 = p_4 = 5$, however without evidence this is not incentive compatible (t_1 and t_2 would imitate t_3 or t_4) and there is a large cost to allocating a_2 to types t_1 and t_2. In fact, we can show that without evidence the optimal mechanism allocates the outside option to types t_1 and t_2 and allocates the mixture $\frac{1}{3}a_1 + \frac{2}{3}a_2$ to types t_3 and t_4 at the price $t_3 = t_4 = 13/3$. The payoff to all types except t_4 is zero, the payoff to type t_4 is $1/3$.

Suppose that we naively attempt to apply the evidence structure used in the proof of Proposition 4. That is $E(t) = \{t_1, t_2\}$ for types t_1 and t_2 and $E(t) = T$ otherwise. We can show that, with this evidence structure, the new optimal mechanism allocates a_1 to t_2 for sure at a price of approximately $p_2 = 0.69$, so that type t_2 is better off. However, the new mechanism allocates the mixture $0.2857a_1 + 0.7143a_2$ to types t_3 and t_4 at a price of $p_3 = p_4 = 4.4286$, yielding a payoff of $0.2857 < 1/3$ to type t_4, so that t_4 is worse off.

Intuitively, since types t_3 and t_4 cannot imitate types t_1 and t_2, it is possible to generate more revenue from types t_1 and t_2 by selling a_1 without the concern that types t_3 and t_4 will deviate. Since t_2 is now attaining a positive payoff the IC $t_2 \to \{t_3, t_4\}$ constraints are relaxed so that the principal can now sell a mixture with a higher probability of a_2 to types t_3 and t_4. However, since the difference between the valuations of types t_3 and t_4 is lower when the probability of a_2 is higher, type t_4's information rent is reduced.

Similarly, we can show computationally that there is no interim Pareto improvement for any deterministic evidence structure.[3]

Note that this example does not satisfy either condition (3a) or condition (3b). For condition (3b), the principal's preferences are type dependent (there is a cost of providing a_2 only for types t_1 and t_2). Note that condition (3a) is violated since type t_1 is willing to pay more than t_4 to move from the outside option to a_2, while conversely type t_4 is willing to pay more than t_1 to move from the outside option to a_1.

[3]This can be done by iterating the relevant linear program over all possible combinations of incentive constraints in MATLAB

4 References

Armstong, Mark (2006). "Recent Developments in the Economics of Price Discrimination", in *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress of the Econometric Society, Volume II*, Blundell, Newey and Persson, eds. Cambridge University Press.

Bergemann, Dirk, Benjamin Brooks and Stephen Morris (2015). "The Limits of Price Discrimination", *The American Economic Review*, forthcoming.

Berheim, B. Douglas, Bezalel Peleg and Michael D. Whinston (1987). "Coalition Proof Nash Equilibria I. Concepts", *Journal of Economic Theory*, 42, 1, pp. 1-12.

Beyer, Anne, Daniel A. Cohen, Thomas Z. Lys and Beverly R. Walther (2010). "The Financial Reporting Environment: Review of the Recent Literature", *Journal of Accounting and Economics*, 50, pp. 296-343.

Board, Simon and Marek Pycia (2014). "Outside Options and the Failure of the Coase Conjecture", *The American Economic Review*, 104, 2, pp. 656-671.

Bicchieri, Cristina and Oliver Schulte (1996). "Common Reasoning about Admissibility", *Erkenntnis*, 45, pp. 299-325.

Bull, Jesse and Joel Watson (2004). "Evidence Disclosure and Verifiability", *Journal of Economic Theory*, 118, 1, pp. 1-31.

Bull, Jesse and Joel Watson (2007). "Hard Evidence and Mechanism Design", *Games and Economic Behavior*, 58, pp. 75-93.

Dye, Ronald A. (1985). "Disclosure of Non-Proprietary Information", *Journal of Accounting Research*, 23, 1, pp. 123-145.

Dye, Ronald A. (1986). " Proprietary and Non-Proprietary Disclosures", *Journal of Business*, 59, pp. 339-366.

Fudenberg, Drew and Jean Tirole (1991). "Perfect Bayesian Equilibrium and Sequential Equilibrium", *Journal of Economic Theory*, 53, pp. 236-260.

Giovannoni, Francesco and Daniel J. Seidmann (2007). "Secrecy, Two-Sided Bias and the Value of Evidence", *Games and Economic Behavior*, 59, pp. 296-315.

Green, Jerry R. and Jean-Jacques Laffont (1986). "Partially Verifiable Information and Mechanism Design", *The Review of Economic Studies*, 53, 3, pp. 447-456.

Grossman, Sanford J. (1981). "The Information Role of Warranties and Private Disclosure about Product Quality", *Journal of Law and Economics*, 24, pp. 461-484.

Grossman, Sanford J. and Oliver D. Hart (1980). "Disclosure Laws and Takeover Bids", *Journal of Finance*, 35, pp. 323-334.

Hagenbach, Jeanne, Frederic Koessler and Eduardo Perez-Richet (2014). "Certifiable Pre-Play Communication: Full Disclosure", *Econometrica*, 83, 3, pp. 1093-1131.

Jovanovic, Boyan (1982). "Truthful Disclosure of Information." *The Bell Journal of Economics*, 13, 1, pp. 36-44.

Jung, Woon-Oh, and Young K. Kwon (1988). "Disclosure When the Market is Unsure of Information Endowment of Managers", *Journal of Accounting Research*, 22, 1, pp. 146-153.

Kamenica, Emir and Matthew Gentzkow (2011). "Bayesian Persuasion", *American Economic Review*, 101, 6, pp. 2590-2615.

Kartik, Navin and Olivier Tercieux (2012). "Implementation with Evidence", *Theoretical Economics*, 7, 2, pp. 323-335.

Lanen, William N. and Robert E. Verrecchia (1987). "Operating Decisions and the Disclosure of Management Accounting Information", *Journal of Accounting Research*, 25, pp. 165-189.

Maskin, Eric (1999). "Nash Equilibrium and Welfare Optimality", *Review of Economic Studies*, 66, 1, pp. 23-38.

McAdams, David (2011). "Discounts for Qualified Buyers Only". Mimeo.

Milgrom, Paul (1981). "Good News and Bad News: Representation Theorems and Applications", *Bell Journal of Economics*, 12, pp. 380-391.

Milgrom, Paul and John Roberts (1896). "Relying on the Information of Interested Parties", *Rand Journal of Economics*, 17, pp. 18-32.

Mussa, Michael and Sherwin Rosen (1978). "Monopoly and Product Quality", *Journal of Economic Theory*, 18, 301-317.

Pae, Suil (2002). "Discretionary Disclosure, Efficiency and Signal Informativeness", *Journal of Accounting and Economics*, 33, pp. 279-311.

Penno, Mark C. (1997). "Information Quality and Voluntary Disclosure", *Accounting Review*, 72, pp. 275-284.

Perez-Richet, Eduardo (2014). "Interim Bayesian Persuasion: First Steps", American Economic Review, 104, 5, pp. 469-474.

Schmalensee, Richard (1981). "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination," *American Economic Review*, 71, 1, pp. 242-247.

Schwartz, Marius (1990). "Third-Degree Price Discrimination and Output: Generalizing and Welfare Result," *American Economic Review*, 80, 5, pp. 1259-1262.

Sher, Itai and Rakesh Vohra (2015). "Price Discrimination Through Communication", *Theoretical Economics*, 10, pp. 597-648.

Stiglitz, Joseph E. (1977). " Nonlinear Pricing and Imperfect Information: The Insurance Market", *Review of Economic Studies*, 44, pp. 407-430.

Tirole, Jean (2015). "From the Bottom of the Barrel to the Cream of the Crop: Sequential Screening with Positive Selection", mimeo.

Varian, Hal R. (1985). "Price Discrimination and Social Welfare," *American Economic Review*, 75, 4, pp. 870-875.

Verrechia, Robert E. (1983). "Discretionary Disclosure", *Journal of Accounting and Economics*, 5, pp. 365-380.

Verrechia, Robert E. (1990). "Endogenous Proprietary Costs Through Firm Interdependence", *Journal of Accounting and Economics*, 12, pp. 245-250.

5 Appendix

5.1 Proof of Lemma 1

The Regularity assumption implies that there is a unique $v_1^* \in (0,1)$. We must have that $v_k^* \in (0, v_{k-1}^*)$, since given the distribution $F(v|v \leq v_k^*)$ both v_k^* and 0 give zero revenue (v_k^* gives zero revenue given our assumption that $F(v)$ is strictly increasing). Therefore, the sequence is strictly decreasing and bounded below by zero. The sequence therefore converges to a limit in $[0, v_1^*)$, by the Monotone Convergence Theorem.

Suppose the limit is strictly greater than 0. Then there exists a $v^* > 0$ such that $lim_{k\to\infty} v_k^* = v^*$.

However, any optimal price given the distribution $F(v|v \leq v^*)$ is $v' < v^*$, since charging v^* will yield zero revenue against the distribution $F(v|v \leq v^*)$. Any optimal price, $v(c)$ on the distribution $F(v|v \leq c)$ solves

$$v(c) \in argmax_{v\in[0,c]} v(1 - F(v|v \leq c)).$$

Given our assumptions, $v(1-F(v \leq c))$ is jointly continuous in (v, c). By the Maximum theorem, $argmax_{v\in[0,c]} v(1-$

$f(v|v \leq c)$ is an upper hemicontinuous correspondence. It follows that

$$v^* = lim_{k\to\infty} v_k^*$$

$$= lim_{c\to v^*} min\{argmax_{v\in[0,c]} v(1 - F(v|v \leq c)\} \in argmax_{v\in[0,v^*]} v(1 - F(v|v \leq v^*).$$

But since $v' < v^*$ for any $v' \in argmax_{v\in[0,v^*]} v(1 - F(v|v \leq v^*)$ this is a contradiction.

Therefore, the limit of $\{v_k^*\}$ as $k \to \infty$ must be zero, as required.

5.2 Proof of Lemma 2

We first note that, for all $t \in \tilde{T}$, $t' \in T \setminus \tilde{T}$, $t < t'$. Suppose not. Then $t > t'$ for some $t \in \tilde{T}$, $t' \in T \setminus \tilde{T}$. But then

$$u(g(t'), t') \geq u((0,0), t')$$

$$\Rightarrow u(g(t'), t) > u((0,0), t) = u(g(t), t),$$

where the strict inequality follows by condition 3 (SCP). Note that whenever $a \neq 0$, $a \succ 0$.

Now, suppose that there exists a triple (ϵ, t, t') with $\epsilon > 0, t \in \tilde{T}$, and t' satisfying $a_{\hat{g}'}(t') \neq 0$, such that

$$(1 - \epsilon)u(\hat{g}(t'), t) + \epsilon u((0, M), t) > u((0,0), t),$$

where M is sufficiently high that $v(\hat{g}(t''), t'') < v((0, M), t'')$ for all $t'' \in T$. Such an M always exists since T is compact and hence the set of outcomes of the mechanism is bounded.

Then, by condition 3 and since $t' > t$ for all $t \in \tilde{T}, t' \in T \setminus \tilde{T}$:

$$(1 - \epsilon)u(\hat{g}(t'''), t''') + (1 - \epsilon)u(0, M, t''') \geq$$

$$(1 - \epsilon)u(\hat{g}(t'), t''') + (1 - \epsilon)u(0, M, t''') >$$

$$u((0,0), t''')$$

for all $t''' \in T \setminus \tilde{T}$. Note that the strict inequality follows by condition 3, since $(1 - \epsilon)a_{\hat{g}}(t) + \epsilon 0 \succ 0$.

It follows that $\hat{g}$ is not optimal, since replacing $\hat{g}(t)$ with $(1-\epsilon)\hat{g}(t) + \epsilon(0, M)$ would satisfy the participation constraints, maintain incentive compatibillity and be preferred by the principal.

Therefore, we must have that for all $(\epsilon, t \in \tilde{T}, t' \in T \setminus \tilde{T})$:

$$(1 - \epsilon)u(\hat{g}(t'), t) + \epsilon u((0, M), t) \leq u((0,0), t)$$

Letting $\epsilon \to 0$ we have

$$u(\hat{g}(t'), t) \leq u((0,0), t),$$

as required.