Learning and Evidence in Principal-Agent Environments*†

Kym Pram‡

University of Nevada, Reno

October 8, 2019

Abstract

I explore the welfare consequences of costly evidence acquisition in a broad class of contracting environments. An initially uninformed agent contracts with a principal. Before choosing whether to participate in a mechanism, the agent can observe, at a cost, a payoff-relevant signal which can be credibly disclosed to the principal. The principal may commit to a mechanism in which allocations are contingent on disclosure of a signal realization. I find that the principal's expected payoff is either non-increasing or U-shaped in the cost of evidence, and derive a condition that precisely distinguishes the two cases. In contrast, the agent's payoff is maximized at intermediate costs of evidence acquisition. Applications include insurance and labor markets,

*Latest version available at `https://sites.google.com/view/kympram`.

†I would like to thank Piero Gottardi, Andrea Galeotti, Alessandro Pavan and Eddie Dekel for their invaluable support and feedback. I would also like to thank Matteo Foschi, Maxim Goryunov, Seetha Menon and a seminar audience at the EUI for their helpful comments.

‡kpram@unr.edu

and public procurement. Further developing the application to insurance markets, I compare the insurer's profit and the agent's welfare between the cases in which evidence can and cannot be contracted upon. I characterize a set of parameter values for which the agent is strictly worse off — and aggregate welfare may be reduced — when evidence can be contracted upon and a set of parameter values for which allowing evidence to be contracted upon induces a Pareto improvement. The results are relevant to the policy debate over the use of genetic testing in health and life insurance markets.

1 Introduction

In many economic environments, agents can acquire information that can be credibly disclosed to another party. For example, medical tests both provide new information to the person taking the test and can be disclosed to an insurance company when purchasing insurance. I will refer to information of this form as endogenous evidence. Is the availability of endogenous evidence good or bad for consumers, for firms, for market efficiency? Existing work on the welfare effects of information considers environments in which information is not endogenous (or its acquisition does not entail learning) or in which information cannot be credibly disclosed. The presence of both of these features critically effects who benefits from information and under what conditions they do so.

The contribution of the paper is twofold. Firstly, I develop a general framework for analyzing the welfare effects of endogenous evidence in adverse selection environments. I use this framework to show how payoffs vary with the cost of evidence. Secondly, I build on these results to analyze a policy question in insurance regulation: should insurance companies be allowed to base the terms of insurance contracts on the result of a genetic test?

In the environments modeled in this paper, information comes in the form of *evidence*: information which can be proved to a third party but which the holder can choose whether or not to disclose. The environments also feature *learning*; an agent who acquires evidence learns about a state of the world that they did not initially know. Evidence is acquired endogenously, at a cost, and the incentives to acquire evidence are shaped by the contract (or menu of contracts) offered by a principal. In a general class of environments with these features I show how the equilibrium payoffs to the principal and the agent depend on the cost of acquiring evidence and on the *value of information*: the benefit to the principal of having information publicly known.

Formally, the model features an initially uninformed agent who contracts with a principal. The agent has access to evidence: a signal that is informative about a payoff-relevant state of the world and can be credibly disclosed to the principal. Before the agent acquires evidence, the principal commits to a contract that specifies an outcome if any specific realization of the signal is disclosed, as well as an outcome if no evidence is disclosed. Outcomes are two-dimensional and can be thought of as an action and a payment. The agent always prefers lower actions and higher payments, while the reverse is true for the principal.

The principal faces a tradeoff between discouraging learning and internalizing the cost of evidence. On the one hand, suppose that the principal specifies a contract that does not induce the agent to acquire evidence. Then the agent has the option to learn, at a cost, whether or not he wants to participate. For example, suppose that the environment is a monopoly goods market and the principal offers a fixed price for the good irrespective of any evidence. The consumer may learn about his value for the good and drop out of the market if his value is low. In order to discourage this, the principal will need to either lower the demanded payment or distort the action away from the ex-ante first best (we will see examples of both cases).

On the other hand, if the principal does induce the agent to acquire evidence he will need to leave the agent rents ex-post. This is because evidence acquisition is costly, therefore the agent will need to be reimbursed for the cost of evidence: otherwise, he would prefer not to participate in the mechanism.

The first effect (learning) becomes more severe as the cost of evidence becomes *lower*, in which case it becomes cheaper to learn. The second effect becomes more severe as the cost of evidence becomes *higher*. Therefore, the principal will require the agent to present evidence when the cost of evidence is below some threshold and will offer an unconditional contract if the cost of evidence is above that threshold.

However, if the cost of evidence is sufficiently high, the optimal contract for the principal is simply to offer the ex-ante optimal action and charge the agent's ex-ante willingness to pay. This implies that above some threshold the optimal contract is constant in the cost, as are payoffs. If this threshold cost is higher than the cost at which the principal is indifferent between inducing evidence acquisition or not doing so, the principal's payoffs are non-monotone (in particular, U-shaped) in evidence. If not, the principal's payoffs are non-increasing in the cost of evidence.

I derive a condition that precisely distinguishes the two cases. If payoffs are quasilinear in payments and the agent's value of the outside option is type-independent, the condition is particularly simple: payoffs are U-shaped in the cost of evidence if and only if the value of information is less than the expected upside risk in the agent's payoff from the ex-ante optimal action. In the general case, the condition implies that the principal's payoffs are U-shaped in the cost of evidence acquisition whenever the value of public information is negative (for example, in insurance markets).

As a consequence, when the value of information is high, all benefits from evidence accrue to the principal. In this case, the principal is best off when the cost of evidence is low. When the value of information is low, the principal is best off when evidence is very costly and the agent is best off with intermediate levels of cost, since at these ranges of cost the principal may leave rents to the agent to discourage learning.

If payoffs are quasilinear in payments, so that aggregate welfare is well defined, aggregate welfare is maximized for low costs of evidence. This is driven by the fact that the value of information is always (weakly) positive when payoffs are quasilinear. For general payoffs, maximization of aggregate welfare is sensitive to the definition of aggregate welfare used.

Conversely, aggregate welfare is smallest for an intermediate cost of evidence in the U-shaped case, and smallest for high costs of evidence in the nonincreasing case; this is also true for general payoffs and any reasonable measure of aggregate welfare.

These comparative statics are significant for the applications in two ways : Firstly, when new forms of evidence are made available by new technologies, a decrease in cost can be seen as a proxy for technological improvements. Therefore, the results tell us who benefits as technologies for producing evidence improve. Secondly, the case in which the principal's payoff is non-monotone in the cost of evidence is precisely the case in which the agent may benefit from the availability of evidence. Therefore, the results tell us when we should expect, based on the primitives, that the availability of evidence might benefit an agent, despite the principal's strong bargaining power in the environments I consider. Variations in cost can be substantial even within a particular class of evidence: for example, genetic tests for hereditary medical conditions can range for under \$100 to over \$2000 (National Institute of Health, cited in Peter et al, 2017).

An existing literature, surveyed in the following section, considers mechanism design environments in which the agent can learn information but cannot credibly disclose it to the principal. My results differ from the results of that literature in three key ways. Firstly, when information cannot be disclosed the agent may be best off with low costs of information acquisition, while in my model the agent is always best off with intermediate costs. Secondly, when information cannot be disclosed, the principal's payoff may be neither nonincreasing nor U-shaped in the cost of evidence acquistion. Moreover, the connection between the shape of the players' equilibrium values and the underlying payoffs is less stark. Thirdly, the literature on non-disclosable information acquisition typically shows that incentives for information acquisition require 'high-powered' menus of contracts, in which the response of the contract offered to the agent's information is greater than first best: this is not the case when the information acquired can be used as evidence.

Examples of economic environments that can be modeled in the framework of this paper include:

1. An insurance company contracts with a consumer who has access to a genetic test. The test provides information about the consumer's risk of a costly disease and can be disclosed to the insurer. In this example the agent is the consumer and the principal is a monopoly insurer (more broadly the results may give some insight into markets which are not monopolistic but where firms have some market power).

2. A local government contracts the provision of a public project to a firm. The firm can commission a technical report which provides information

about the cost and feasibility of the project. This report can be shown to the local government and the terms of the contract may depend on the information contained in the report. In this example the local government is the principal and the firm is the agent.[1]

3. A potential employee completes an educational course before seeking a job with a monopsonistic employer. While completing the course the potential employee learns his own ability, which affects his expected outside option. The qualification and grades received can be shown to the employer and the terms (or existence) of an employment contract may depend on these. In this example the employer is the principal and the potential employee is the agent.[2]

In the second part of the paper I develop the genetic testing example in more detail. In particular, I compare two ways in which an insurance market might be regulated: (i) the insurer is allowed to ask the consumer for the results of a genetic test and condition the terms of an insurance contract on the results; (ii) the insurer must offer the same contract (or menu of contracts) to the consumer regardless of genetic test results.

I show that if the cost of the test is low and the ex-ante probability of low-risk test results is high, or if the cost lies in an intermediate range, then the consumer is strictly worse off in case (i). In contrast, if the cost of the test is low and the ex-ante probability of a high-risk test result is high, then the consumer is no worse of under case (i) compared to case (ii), while the insurer is strictly better off. Therefore, under the latter set of parameter

[1] Of course, the contracter will eventually learn the cost of production if the project is undertaken. However, this is compatible with my model if, for example, final production costs will have to be learned after the contract is signed, even if a preliminary estimate has been made before contracting.

[2] This model of education abstracts from its role in human capital formation and as a signal of exogenous private information. While those effects are important, education as endogenous evidence illustrates another channel through which it may affect labor markets.

values, allowing insurance companies to condition on the results of a genetic test is Pareto improving.

Moreover, it is possible that aggregate welfare (measured as a weighted sum of the insurer's and consumer's payoff) is lower when evidence can be contracted upon than when it cannot. This can occur when costs are in an intermediate range so that evidence will be acquired in the optimal mechanism when it can be contracted upon but not acquired when it cannot be contracted upon. Aggregate welfare is lost through both the direct cost of evidence acquisition, and through the Hirshleifer effect – the observation that ex-ante insurance is more socially valuable than ex-post insurance when the consumer is risk-averse. For some parameter values this loss is not offset by the elimination of the adverse selection problem when evidence can be used in the contract.

Regulation of the use of genetic information in insurance contracting is the subject of an ongoing policy debate. In 2008 the US congress passed the Genetic Information Nondiscrimination Act (GINA). Among other provisions, the GINA prohibits *health* insurers from denying coverage, or altering the contracts offered, based on the result of a genetic test. But the GINA does not extend to life or disability insurance. Many interested parties, including medical groups, have voiced concerns that this restriction in scope may harm consumers, leading to ruinously high premiums or denial of coverage for individuals with a predisposition to costly health issues (Erwin, 2008). On the other hand, industry players have voiced concerns that extending the GINA to life insurance could lead to unravelling in the market. A statement by the American Academy of Actuaries argues that barring companies from contracting on test results in this market "could lead to adverse selection and impact the stability of rates."[3]

[3]https://www.nytimes.com/2014/04/08/science/fearing-punishment-for-bad-

My results illustrate how the validity of these positions depends on the specific parameters of the market. As the costs of genetic testing fall, extending the GINA to life insurance could benefit consumers if low risk types, relative to the ex-ante expectation, are more likely. Conversely, extending the GINA to life insurance could (weakly) harm both sides of the market if high risk types are relatively more likely. Which is the case is a question for empirical research.

1.1 Literature Review

A number of existing theoretical papers model insurance markets with genetic testing. The most closely related is Doherty and Thistle (1996) which considers genetic information in the context of a perfectly competitive — rather than monopolistic — market. Although their focus is on a different question — why and under what conditions consumers choose to become informed about their risk — their results imply that consumers are better off when evidence can be contracted upon in the competitive case (given that an equilibrium exists). In their model, with some probability the consumer has evidence available exogenously, a feature not present in my model.

In less directly comparable models, Hoy and Polborn (2000) consider a competitive market in which insurers cannot screen consumers using the level of insurance, and Strohmenger and Wambach (2000) consider a model in which the cost of treatment may exceed consumers' ability to pay. Hoel et al (2006) consider a model in which consumers have intrinsic preferences over when uncertainty is resolved, separately from the instrumental value of information. Barigozzi and Henriet (2011) consider a model in which genetic information has instrumental value in treatment or prevention decisions, where preven-

genes.html

tion reduces the size of loss. Peter et al (2017) consider a similar model in which prevention reduces the probability of a loss (and is unobservable). Both of the two latter papers compare four policy regimes and find that a duty to disclose results of a test Pareto dominates all alternatives, while an information ban (such as the GINA) is Pareto dominated by all other alternatives.

In contrast to these papers I consider a monopolistic insurance market. As some empirical papers (Chiappori et al, 2006; Cohen and Einav, 2007; Dafny, 2010) find evidence of market power in insurance markets, monopolistic models are empirically relevant.

Lagerlof and Schottmuller (2016a, 2016b) consider a monopolistic model of insurance with genetic testing which features learning but not evidence: that is, the consumer learns about his risk type by taking a test but cannot prove what he knows to the insurer. Their model can be seen as a benchmark against which we can compare outcomes when information can be credibly disclosed and contracted upon.

Returning to the general model, there is a large literature examining mechanism design with endogenous information and a sizeable literature considering mechanism design with evidence. Here I emphasize the contribution of my paper over the most closely related existing work. McAdams (2011), in the specific context of a unit-valuation monopoly model analyzes the impact of costly evidence acquisition without learning. He shows that aggregate welfare is U-shaped in the cost of evidence while the monopolist's profit is non-increasing. Some of my results can be seen as a generalization of a special case of McAdams' model: in which the buyer holds no initial private information.

More broadly, a sizable literature characterizes optimal mechanisms with exogenous evidence: that is, when an agent (or set of agents) can prove some information to the principal, but the available evidence is costless or exogenous to the chosen mechanism. These environments feature *evidence* but not learning. Sher and Vohra (2015) characterize optimal mechanisms in a unit valuation monopoly model with arbitrary evidence structures. Ben-Porath, Dekel and Lipman (2017) characterize optimal mechanisms with evidence in a class of allocation problems without transfers. Glazer and Rubinstein (2004, 2006) consider a particular class of environments where a principal chooses whether to accept or reject a request and, as well as characterizing the optimal mechanism, consider the question of whether commitment is necessary. Another strand of literature, following Green and Laffont (1986) investigates general issues relating to partial or full implementation with evidence.

Another, larger strand of literature characterizes optimal mechanisms in environments with endogenous information that cannot be credibly disclosed. These environments feature *learning* but not *evidence*. Early examples include Barzel (1977), Craswell (1988) and Demski and Sappington (1987). A more recent strand follows papers by Cremer and Kahlil (1992) and Cremer et al (1998a; 1998b) in the procurement context. A more recent example, generalizing the work of Cremer et al. is Szalay (2009).

The most closely related of these papers are Cremer et al (1998a) and Szalay (2009). A key result of these papers is that contracts that induce information of non verifiable information feature 'high-powered' incentives. That is, the contract offered to low types specifies an action that is further below the mean action than in the optimal contract with exogenous information, while the contract offered to high types is further above the mean. While the model I analyze is too general for 'low type' and 'high type' to be well defined (I do not impose a sorting condition such as single crossing), two examples that I analyse suggest that the opposite finding is often true with endogenous *evidence*: when the principal chooses to induce evidence acquisition the actions specified ex-post are less variable than in case of either exogenous non verifiable information or exogenous evidence.

My welfare results also differ from the results of this literature. In the case of endogenous non verifiable information, the principal's payoff need not be either nonincreasing or U-shaped in the cost of evidence and the agent's payoff may be maximized at low costs, rather than intermediate costs (see Hoppe and Schmitz, 2010, for explicit examples of both possibilities).

A further stand of literature characterizes optimal selling mechanisms and optimal auctions with endogenous information. Examples include Persico (2000) and Shi (2012). Other papers, such as Bergemann and Valimaki (2002), consider the implementation of efficient allocations with endogenous information.

Aside from the feature that the information acquired in these models cannot be credibly disclosed, these papers consider more specific settings while the present paper considers a very general class of adverse selection environments. On the other hand, these papers tend to consider environments in which information is chosen from a large set of possible information structures. In contrast, in my model evidence acquisition is an all-or-nothing choice. Neither approach is a generalization of the other, since these models generally restrict the relation between information choice and the cost of information acquisition to follow a particular form. Modeling information acquisition as an all-or-nothing choice is relevant in many cases where the information acquisition technology is given exogenously, for example with a genetic test.

De Marzo, Kremer and Skrzypacz (2015) consider a setting that does involve both evidence and learning, however the model and focus are different in a number of ways. In their model, a seller can acquire a signal (test result) and disclose it to a market. Since disclosure is to a competitive market of buyers, rather than a principal, the market reacts in a sequentially rational manner to any disclosure rather than shaping incentives to acquire evidence through commitment to a mechanism. The market does not observe which test, from a set of available tests, the seller has taken. Their main result a characterization of which test the seller chooses in equilibrium: the signal that minimizes the expected price conditional on non-disclosure (tests in this model may produce a null result).

The literature on certification (such as Lizzeri, 1999) studies similar issues. This literature does not feature learning on the part of the player whose type is being certified, and typically features a certification intermediary who provides evidence of the (exogenous) type and disclosure to a market rather than a principal.

A number of papers investigate the welfare consequences of *exogenous* hard information in mechanism design environments. Pram (2017) provides sufficient conditions under which *some* exogenous evidence structure can induce an interim Pareto improvement. Schmalensee (1981), Varian (1985) and Schwartz (1990), in increasing generality, find conditions on demand under which third degree price discrimination improves aggregate welfare in monopoly markets. Bergemann, Brooks and Morris (2015) characterize the set of consumer surplus - profit pairs that are feasible across all (exogenous) information structures for a monopolist facing a unit valuation buyer. Roesler and Szentes (2017) characterize optimal (costless) learning (of information that cannot be disclosed) in a monopoly model. All of these papers differ from the present paper in significant ways. In particular, none of these papers allow for both learning and evidence.

2 General Results

2.1 Model

In the general model, an initially uninformed agent contracts with a principal. The agent has access to an evidence-acquisition technology: at cost c, the agent can learn his payoff relevant type and can disclose this to the principal. For example, in the insurance application c would be the cost of taking a genetic test. The principal commits to a mechanism before the agent decides whether to acquire evidence. The mechanism may specify different outcomes contingent on the type disclosed, and contingent on no disclosure. I make very few assumptions about the payoffs or information structure, except that allocations are two dimensional (for example, a price and a level of trade) and both the principal's and agent's payoffs satisfy natural monotonicity properties in each dimension (for example, the agent prefers a lower price and a higher level of trade while the reverse is true for the principal).

Formally, a principal (she) faces an agent (he) with an unknown type $t \in T$, where T is a finite set. Initially, both the principal and the agent know only the prior, $\pi \in \Delta(T)$, where $\Delta(T)$ denotes the set of probability measures on T. The agent has access to evidence: at cost c the agent can learn t and if he does so he can credibly disclose this information to the principal.

Outcomes are denoted by a pair $(a, p) \in A \times \mathbb{R}$ where A is a compact, Borel, subset of $\mathbb{R}$. Let $\Delta(A \times P)$ be the set of probability measures on $A \times P$ equipped with the restriction of the Lebesgue measure on $\mathbb{R}^2$ to $A \times P$. The agent has access to an outside option, denoted $0 \in A \times P$. Different types of the agent may value the outside option differently

Ex-post payoffs are given by:

$$u : A \times P \times T \to \mathbb{R}$$

for the agent and

$$v : A \times P \times T \to \mathbb{R}$$

for the principal. Preferences over lotteries satisfy the expected utility hypothesis. With some abuse of notation, let $u(\phi, t)$ and $v(\phi, t)$ denote $E_\phi[u((a, p), t)]$ and $E_\phi[v((a, p), t)]$, respectively, for $\phi \in \Delta(A \times P)$. We make the following assumptions on preferences:

Assumption 1 *For each $t \in T$, $u(\cdot, \cdot, t)$ is weakly increasing in a and strictly decreasing in p.*

Assumption 2 *For each $t \in T$, $v(\cdot, \cdot, t)$ is weakly decreasing in a and strictly increasing in p.*

Assumption 3 *For each $t \in T$, $u(\cdot, \cdot, t)$ and $v(\cdot, \cdot, t)$ are jointly continuous in a and p.*

Assumption 4 *Let $\phi \in \Delta(A \times P)$ and let $\bar{p} = E_\phi(p)$ and $\hat{a} = marg_A\phi$. Then $v(\hat{a}, \bar{p}, t) \geq E_\phi v((a, p), t)$ and $u(\hat{a}, \bar{p}, t) \geq E_\phi u((a, p), t)$ for all $t \in T$.*

Assumption 5 *For each $t \in T, \alpha \in \Delta(A), c \in \mathbb{R}$ there exists p such that:*

$$u(\alpha, p, t) = c.$$

Assumptions 1 and 2 are monotonicity requirements.[4] Assumption 3 is a standard continuity assumption that is necessary to guarantee existence of

[4]Note that, in a given application, if it is natural for a consumer to prefer higher actions, a', and a principal to prefer higher payments, p', we can simply redefine $a = -a'$ and $p = -p'$.

an optimal mechanism. Assumption 4 requires that both the principal and agent are weakly risk averse over payments for any given state $t \in T$, and imposes a weak form of separability in preferences over actions and payments. Assumption 5 ensures that the principal cannot extract an unbounded payment from the agent. All five assumptions are commonly satisfied in applications.

Before any evidence is acquired, the principal commits to a mechanism. A mechanism specifies an outcome function $g : T \cup \{N\} \to \Delta(A \times P)$. The interpretation is that the mechanism specifies an outcome for each signal realization, $t \in T$ that might be presented to the principal, as well as an outcome, $g(N)$, if the agent presents no evidence. A seemingly more general mechanism would also allow the mechanism to be conditioned on cheap-talk messages in the event that the agent acquires the signal but does not present it. In the Appendix I show that without loss of generality the agent presents all acquired evidence to the principal.

The agent incurs a separable cost, c, if he chooses to acquire evidence, so that the agent's ex-ante payoff when evidence is acquired on-path is:

$$E_t[u(g(t), t)] - c,$$

noting that $g(t)$ may be a lottery over outcomes.

The timing is as follows:

Period 1: The principal commits to a mechanism	Period 2: The agent chooses whether to acquire evidence	Period 3: Evidence is presented and the outcome is realized

Figure 1: Timing of the problem

The agent can choose to take the outside option either before or after acquiring evidence. Preventing the agent from opting out before acquiring evidence imposes an ex-ante individual rationality constraint. Preventing the agent from opting out after acquiring evidence imposes both a moral hazard constraint and an ex-post individual rationality constraint.

The ex-post individual rationality constraint is relevant when evidence is acquired on-path: since the agent knows his type after acquiring evidence he must prefer to accept the contract offered ex-post. The moral hazard constraint is relevant when evidence is not acquired on path. If the agent chooses to acquire evidence off-path, he may prefer to take the outside option contingent on certain types. Therefore, if the principal wishes to prevent the agent from acquiring evidence (which is an unobservable decision), the agent must not expect to do too much better, ex-post, by choosing to take the outside option contingent on learning that the type is in some set. This is the moral hazard constraint

Note that evidence acquisition is a once and for all decision, as is common in the literature.

Let $d \in \{Y, N\}$ with $d = Y$ if the agent acquires evidence and $d = N$ if the agent does not acquire evidence. An optimal mechanism for the principal solves the following maximization problem:

$$max_{(g(t))_{t \in T}, g(N), d} \mathbf{1}\{d = Y\} E_t v(g(t), t)) + \mathbf{1}\{d = N\} E_t v(g(N), t)), \quad s.t.$$

$$E_t u(g(N), t) \geq E_t[u(0, t)]. \quad (IR_N) \tag{1}$$

$$E_t u(g(N),t) \geq \sum_{t \in T} E_t max\{u(g(N),t), u(0,t)\} - c. \ \ (MH) \tag{2}$$

$$E_t[u(g(t),t)] \geq E_t[u(0,t)] - c. \ \ (IR_Y) \tag{3}$$

$$u(g(t),t) \geq u(0,t) \ \forall t \in T. \ \ (EPIC) \tag{4}$$

$$d = Y \Leftrightarrow E_t[u(g(t),t)] - c \geq E_t[u(g(N),t)] \ \ (O) \tag{5}$$

The interpretation of the constraints is as follows: (1) ensures that the agent prefers to participate in the mechanism when evidence is not acquired on path, (2) ensures that if evidence is not acquired on path then the agent does not prefer to acquire evidence and then participate in the mechanism only conditional on certain signal realizations, (3) ensures that the agent prefers to participate in the mechanism when evidence is acquired on-path, (4) ensures that, after observing each signal realization, the agent prefers to present the signal and participate in the mechanism than to take the outside option, (5) ensures that the agent prefers to follow the principal's recommendation on whether or not to acquire evidence.

Under assumptions 1-5 an optimal mechanism is guaranteed to exist. Moreover it features no randomization in the agent's evidence acquisition decision:

Lemma 1 *Under Assumptions 1-5 an optimal mechanism exists. Moreover, without loss of optimality the agent's evidence acquisition decision is deterministic.*

Suppose the agent randomizes over evidence acquisition. Then the principal must be indifferent between the contract chosen when the agent has evidence and the contract chosen when the agent does not, otherwise the principal is better off altering one or the other slightly to break the indifference. But then it is optimal to simply offer one or the other. Moreover, a contract in which the agent randomizes over the evidence acquisition decision is typically strictly worse than one in which the agent acquires evidence for sure. This is because, when there is some probability that the agent does not acquire evidence, the principal will typically offer a contract other than the outside option conditional on no disclosure. This contract may be preferred to the outside option by some ex-post types. This introduces a stronger incentive constraint which limits what the principal can offer conditional on the agent presenting evidence.

A proof of the lemma is contained in the appendix.

2.2 Examples

A number of economic environments can be modeled within this framework. Some examples include:

- A risk-averse consumer purchases insurance from a monopoly insurer. The consumer faces a risk of a loss $\phi \in \{\phi_L, \phi_H\}$ and initially knows only the prior. The consumer can, at a cost, c, obtain a genetic test. The result of the test reveals the risk and can be credibly disclosed to the insurer. In this example, $a \in A$ is the level of insurance (the proportion of the loss that is repaid in the event of a loss), and p is the premium. The state space is $T = \{\phi_L, \phi_H\}$. This example is analyzed in greater detail in Section 3.

- A local government hires a firm to build a public project. The cost of production is unknown and lies a the set T. The firm can commission a technical report to learn the cost and can disclose the results of the technical report to the local government. In this example the firm is the agent and the local government is the principal. The action space is $A = [0, \bar{a}]$, where a is the size of the project and $\bar{a}$ is a maximum desirable size. The payoffs are:

$$v(a, p, t) = V(a) - p.$$

$$u(a, p, t) = p - a \cdot t.$$

Where $V(a)$ is increasing in a. Note that in this example the agent prefers *higher* payments and *lower* project size, while the reverse is true for the principal. However, the environment can be recast in my framework by defining $a' = -a$ and $p' = -p$ and writing payoffs as a function of (a', p').

The cost of evidence acquisition, c, should be interpreted as the difference between the cost of commissioning a preliminary technical report as well as learning final costs when production is undertaken, and the cost of simply learning final costs once the contracter has already commited to producing the project.

- A potential employee may be hired by a monopsonistic employer. The employee may be a low or high productivity type, $t \in \{t_L, t_H\}$, but does not initially know his type. The employer chooses whether to offer a job, $a \in \{0, 1\}$, and at what wage, p. Before interacting with the employer the employee can complete a qualification at some cost, c, which, unlike in the Spence signaling model, is independent of the type. Completing a qualification informs the employee about his type and the qualification, or grade received can be disclosed to the employer. The type is also informative about the employee's expected wage if he enters a different industry, $o(t)$. Payoffs are given by:

$$v(a, p, t) = t \cdot a - p.$$

$$u(a, p, t) = p.$$

While the value of the outside option to the agent is $o(t)$ conditional on type t.

2.3 Results

In this section I show how the cost of evidence determines the equilibrium payoffs to the principal and agent. The main result is that the principal's ex-ante payoff is either U-shaped or monotone non-decreasing in the cost of evidence. In the first case it is highest for extreme costs and first decreases then increases as costs increase up to a point at which it becomes constant. In contrast the agent's payoff is (weakly) highest for intermediate costs.

In the second case (where the principal's payoff is non-increasing), the principal's profit is strictly decreasing in cost up to a point at which it becomes constant. In this case, the agent's ex-ante payoff is always equal to the payoff he receives under the outside option. The second case occurs if and only if the value to the principal of having the type commonly known is sufficiently high, in particular it must be non-negative.

The condition distinguishing the two cases is particularly clear in environments where payoffs are quasilinear and the value of the outside option is type independent. In these environments, the U-shaped case occurs if and only if the value of information is less than the expected upside risk in the value of the ex-ante optimal action to the agent. The expected upside risk can be seen as a measure of the variation in the agent's preferences across types and is closely related to the notion of 'information sensitivity' used in the finance literature (Dang et al, 2015).

A decrease in the cost of evidence can be interpreted as an improvement in the underlying technology used to produce evidence. At an extreme, if the cost is very high this is equivalent to the evidence being unavailable. The results suggest that an improvement in the evidence production technology (in the form of lower costs) may benefit the agent only if the value of information is sufficiently low and the agent's preferences are not too variable across types. The results are driven by the following lemma:

Lemma 2 *There exist threshold costs $c^* > 0$ and $c^{**} \geq c^*$ such that evidence is acquired if on-path if and only if $c < c^*$ while if $c > c^{**}$ then IR_N is binding and MH is non-binding.*

Intuitively, the value to the principal of the best mechanism in which evidence acquisition is discouraged on-path is non-decreasing in cost. Conversely, the value to the principal of the best mechanism in which evidence is acquired on path is strictly decreasing in cost.

First consider the case in which evidence acquisition is discouraged. The key constraint in this case is that the agent must be prevented from learning his type and choosing not to participate in the mechanism if the type is unfavorable (MH). This constraint is relaxed as the cost of evidence acquisition increases. If cost is sufficiently high, this constraint is not binding at the ex-ante optimal contract so that above that point the principal can offer the ex-ante optimal contract with IR_N binding, and MH slack.

Next consider the case in which evidence is acquired. The value to the principal of the best mechanism in which evidence is presented on path is decreasing in cost. In particular, the IR_Y constraint (which becomes stricter as cost increases) is binding. We can show that evidence does better than no evidence when $c = 0$, while the converse is true when c is very high. After establishing the continuity of each value function it follows that a threshold value, c^*, as in the lemma exists.

As noted above, when evidence is not used on path, there is a threshold, $\tilde{c}$, above which the learning constraint (MH) does not bind. Above this value the optimal contract does not depend on cost; below it, the value to the principal is strictly increasing. Letting $c^{**} = max\{\tilde{c}, c^*\}$ establishes the second part of Lemma 2. The principal's payoff will be $U-shaped$ in cost if $c^* < c^{**}$ while it will be non-increasing in cost if $c^* = c^{**}$.

A formal proof of Lemma 2 is contained in the Appendix.

The following Lemma will be useful in the proofs of many following results:

Lemma 3 *If evidence is acquired on-path then IR_Y is binding. If evidence is not acquired on-path then either IR_N or MH is binding.*

Proof: Suppose that evidence is acquired on-path and IR_Y is not binding, that is:

$$E_t[u(g(t),t)] > E_t[u(0,t)] + c.$$

Since $E_t[u(g(t),t)] > E_t[u(0,t)]$ we must have that EPIR is not binding for some $t \in T$. I will show by a variational argument that the principal can do better.

Let $g'(t) = (1-\epsilon)g(t) + \epsilon\hat{g}(t)$ for some $\hat{g}(t) \in A \times P$ such that $v(\hat{g}(t),t) > v(g(t),t)$ and ϵ sufficiently small. Since the principal's payoff is strictly increasing in p, such a $\hat{g}(t)$ always exists. Then it is straightforward to see that all constraints are still satisfied and the principal's payoff is improved.

Now suppose that evidence is not acquired on-path and neither IR_N nor MH are binding in an optimal contract, $g(N)$ — recalling that without loss of optimality $g(t) = 0$ for all $t \in T$ in this case.

A similar variational argument shows that the contract $g(N)$ cannot be optimal. Let $g'(N) = (1 - \epsilon)g(N) + (1 - \epsilon)\hat{g}(N)$ for some $\hat{g}(N)$ such that $E_t[v(\hat{g}(N), t)] > E_t[v(g(N), t)]$ and ϵ sufficiently small. Since IR_N and MH are the only relevant constraints in this case and both are slack under g, the new contract is feasible and improves the principal's payoff, contradicting the optimality of g. ■

Lemmas 2 and 3 imply that the principal's profit is either non-increasing or U-shaped in the cost of evidence. Below the cutoff, c^*, the principal prefers to induce evidence acquisition. In this range of costs, the principal's profit is decreasing in cost because IR_Y is binding and becomes stricter as costs increase. If the interval $(c^*.\tilde{c})$ is nonempty then the principal's profit is increasing in the cost of evidence acqusition. This is because the relevant constraint in this region is MH, which becomes more slack as cost increases (that is, it is easier to discourage learning when the cost of learning is higher). In this intermediate range the agent may obtain positive rents.

Above $\tilde{c}$, the ex-ante optimal contract — the contract that the principal would specify if evidence was not available — is feasible. If the range $(\tilde{c}, c^*)$ is empty then the principal either induces evidence acquisition (for low costs) or specifies the ex-ante optimal contract. Therefore, in this case, the principal's payoff is decreasing up to a point, at which is becomes constant.

The main result of this section provides a condition that distinguishes between the two cases (U-shaped and monotone non-decreasing) based on the underlying payoffs. I first derive the condition for quasilinear environments, where the result is sharper. Analyzing the quasilinear case will also help to build intuition for the general case.

2.3.1 Quasilinear payoffs

In many economic environments preferences are well represented by quasilinear payoffs. Quasilinear payoffs simplify the analysis of mechanisms with endogenous evidence: the principal's payoff from an optimal mechanism in which evidence is acquired is simply the expected social value of the ex-post first best action less the cost of evidence. Moreover, the optimal mechanism simply specifies that the ex-post optimal action is taken following any given disclosure and uses the price to 'reimburse' the cost of evidence acquisition to the agent.

Formally, in this subsection we specialize the ex-post payoffs to:

$$v(a,p,t) = V(a,t) + p.$$

$$u(a,p,t) = U(a,t) - p.$$

Where $V(a,t)$ is decreasing in a for all $t \in T$ and $U(a,t)$ is increasing in a for all $t \in T$.

Let a^* maximize the ex-ante expected surplus given a type-independent action. That is, a^* maximizes:

$$E_t[V(a,t) + U(a,t)]$$

and let a_t^* maximize the ex-post expected surplus for type t. That is, a_t^*

maximizes:

$$V(a_t^*, t) + U(a_t^*, t).$$

With quasilinear payoffs the optimal choice of action, from the point of view of aggregate surplus, is independent of the price charged. When evidence is acquired on-path, the principal does not have to pay rents to distinguish types since failure to disclose a type $\tilde{t}$, will imply that the true type is not $\tilde{t}$ (if no evidence is disclosed then, as we have seen, it is without loss of generality to implement the outside option). The principal is therefore best off specifying the optimal ex-post action and a price that extracts the agent's surplus, following any disclosure. To capture this idea, let $\hat{p}_t(a)$ solve $V(a,t) - U(0,t) = \hat{p}_t(a)$, for an action $a \in A$.

Similarly, if the cost of evidence is sufficiently high that the agent will never acquire evidence, the optimal mechanism will select the ex-ante optimal action and a price that fully extracts the agent's surplus ex-ante. To capture this idea let $\hat{p}(a) = E_t[V(a,t)] - E_t[U(0,t)]$ for an action $a \in A$.

The quasilinear structure then allows a particularly simple characterization of optimal mechanisms when evidence is acquired on-path:

Lemma 4 *If evidence is acquired on-path the optimal mechanism is given by*

$$g(t) = (a_t^*, \hat{p}_t(a_t^*) - c) \;\; \forall t \in T$$

and $g(N)$ is the outside option.

Note that the cost of evidence is reimbursed uniformly across types (in fact, there are other ways of distributing the cost across types, but all lead to the same ex-ante payoffs for the principal and agent). Note that if payoffs are not quasilinear, other ways of distributing the cost across types may be better: for example, if the agent is risk-averse in payments and the principal is risk-neutral, the principal may do better by 'smoothing' payments across types.

An important object in the results will be the *value of information*. The value of information is defined to be the difference between joint surplus when the optimal action is chosen ex-post and when the optimal action is chosen ex-ante:

$$VI = E_t[V(a_t^*, t) + U(a_t^*, t)] - E_t[V(a^*, t) + U(a^*, t)].$$

In the quasilinear case the value of information is always non-negative. As we shall see, this is not true in the general case.

Given Lemma 4, we can derive the following sharp characterization of the cases in which the principal's expected payoff is U-shaped in the cost of evidence:

Proposition 1 *The principal's expected payoff is U-shaped in c if and only if*

$$VI < E_t[max\{U(0,t) - E_t[U(0,t)], U(a^*, t) - E_t[U(a^*, t)]\}].$$

Otherwise it is monotone non-deceasing in the cost of evidence. In both cases the payoff is constant above c^{**}

Suppose that the value of the outside option is type-independent. Then Proposition 1 says that the principal's payoff is U-shaped if and only if the value of information is less than the expected upside risk in the agent's preferences over the ex-ante optimal action, a^*. When the value of the outside option is type-dependent then this condition generalizes to the requirement that the maximum of the difference between the actual and expected values of (i) the ex-ante optimal action and (ii) the outside option is larger than the value of information. Informally, the condition requires that the value of information is not too high and the agent's preferences vary sufficiently across types.

Note that this condition can hold even if the value of information is strictly positive. That is, even though the principal strictly prefers to be informed and can extract the net social surplus when evidence is acquired, decreasing the cost of evidence makes the principal worse off over some interval.

Proposition 1 is proved in two steps. Firstly, I show that the principal's profit is either monotone non-increasing or U-shaped in the cost of evidence. This follows from Lemmas 2 and 3, after showing that the MH constraint must be binding whenever the cost of evidence is sufficiently high that the principal prefers to discourage evidence acquisition, but not so high that the ex-ante optimal contract is feasible.

Secondly, I derive the inequality given in proposition 1. The U-shaped case will occur if the value of the best mechanism which discourages evidence acquisition is higher than the value of the best mechanism with on-path evidence acquisition, at the cutoff value, $\tilde{c}$, at which MH is satisfied with equality given the ex-ante optimal mechanism. The left hand side of the inequality relates to the difference in the principal's payoff from the two mechanisms, before reimbursing the cost of evidence acquisition to the agent. The right hand side is the cost, $\tilde{c}$, which is defined by an 'indifference to learning' condition.

To prove Proposition 1 formally, let c^* denote the cutoff below which the principal prefers to induce evidence acquisition and let $\tilde{c}$ be the cutoff above which $(a^*, E_t[u(a^*, t)] - E_t[u(0, t)])$ is feasible. Suppose that $c^* < \tilde{c}$. Then between c^* and $\tilde{c}$ we can show that MH is binding. Suppose that it is not, then consider the contract $(1-\epsilon)g(N)+\epsilon(a^*, E_t[U(a^*,t)]-E_t[U(0,t)])$ for small ϵ. Since the elements of the mixture satisfy IR_N, this constraint is satisfied by the mixture. Since MH is slack at $g(N)$, is it satisfied for the new contract if ϵ is sufficiently small. However, the principal's payoff is improved, contradicting the optimality of g.

Since MH is relaxed as c increases, while IR_N is unchanged, it is easy to see that the principal's profit is increasing in c over this range. Conversely, for $c<c^*$ IR_Y is binding. Since IR_Y becomes harder to satisfy as c increases, the principal's payoff is decreasing in c over $[0,c^*)$.

It follows that the principal's payoff is non-increasing in c if $c^* \geq \tilde{c}$ and U-shaped in cost if $c<\tilde{c}$. Note that, since the principal's payoff from a mechanism in which evidence is not acquired on-path is increasing in c below $\tilde{c}$ and constant in c above $\tilde{c}$, a necessary and sufficient condition for the principal's payoff to be non-increasing in cost is that the principal does better by inducing evidence acquisition when the cost of acquisition is $\tilde{c}$.

The cutoff $\tilde{c}$ is defined by:

$$E_t[U(0,t)] = E_t[max\{U(0,t), U(a^*,t)-E_t[U(a^*,t)]+E_t[U(0,t)\}]-\tilde{c}.$$

The left hand side is the agent's payoff from $(a^*, E_t[U(a^*,t)]-E_t[U(0,t)])$. The right hand side is the agent's payoff if that same contract is offered and the agent chooses to learn his type and take the outside option when it is preferable.

If the contract offered is $(a^*, E_t[U(a^*,t)]-E_t[U(0,t)])$ then MH is binding when $c=\tilde{c}$, slack if $c>\tilde{c}$ and not satisfied if $c<\tilde{c}$. Notice that by construction $(a^*, E_t[U(a^*,t)]-E_t[U(0,t)])$ is the best contract for the principal that satisfies IR_N. Therefore if evidence is not acquired on-path and $c \geq \tilde{c}$ the optimal contract must be $(a^*, E_t[U(a^*, t)] - E_t[U(0, t)])$.

Re-arranging the implicit expression for $\tilde{c}$ we obtain:

$$\tilde{c} = E_t[max\{U(0, t) - E_t[U(0, t)], U(a^*, t) - E_t[U(a^*, t)]\}].$$

We now compare the best mechanisms with and without on-path evidence acquisition at the cost level $\tilde{c}$. Using the result of Lemma 4, the principal prefers to induce evidence acquisition at $\tilde{c}$ if and only if:

$$\begin{aligned} E_t[V(a_t^*, t) + U(a_t^*, t)] - E_t[max\{U(0, t) - E_t[U(0, t)], U(a^*, t) - E_t[U(a^*, t)]]\}] \\ \geq E_t[V(a^*, t) + U(a^*, t)] \\ \Leftrightarrow VI \geq E_t[max\{U(0, t) - E_t[U(0, t)], U(a^*, t) - E_t[U(a^*, t)]\}]. \end{aligned}$$

Noting that $E_t[V(a^*, t) + U(a^*, t)] - E_t[U(0, t)]$ is the payoff to the principal from $(a^*, \hat{p}(a^*))$, which is feasible when $c = \tilde{c}$. This proves the proposition.

Moreover, it is clear that the agent's payoff is always highest for the intermediate range of costs, $c \in (c^*, \tilde{c})$, when this range is non-empty. In particular, the agent's ex-ante incentive constraint is always binding outside this range. As a convention, in the following proposition I assume that the principal chooses to induce evidence acquistion if he is indifferent to doing to:

Proposition 2 *The agent's ex-ante payoff may be greater than the reservation payoff only if $(c^*, \tilde{c})$ is nonempty and $c \in (c^*, \tilde{c})$.*

As we have seen, when $c > \tilde{c}$, no evidence is acquired and the optimal contract is $(a^*, \hat{p}(a^*))$, at which the ex-ante IR constraint is binding. Similarly, at $c \leq c^*$ evidence is acquired on-path and the ex-ante IR constraint (net of costs) is binding.

As a corollary, we can show which levels of cost minimize and maximize aggregate welfare: defined as the sum of the principal's and the agent's payoff.[5] In the U-shaped case, aggregate welfare is clearly minimized at $c = c^*$. At this level of cost the principal's payoff is minimized and the agent's payoff kept to the outside option.

Aggregate welfare is maximized at $c = 0$: recall that the action specified by the mechanism when $c \geq \tilde{c}$ is a^*, which maximizes aggregate welfare subject to the IR_N constraint. Since this constraint must be satisfied for any $c \in (c^*, \tilde{c}]$ it follows that the mechanism chosen when $c \geq \tilde{c}$ maximizes aggregate welfare across all c for which evidence is not acquired on-path.

Similarly, when $c \leq c^*$ the vector of actions specified by the mechanism is $(a_t^*)_{t \in T}$, independently of the cost. In this region, aggregate welfare from the optimal mechanism will be $E_t[v(a_t^*, t)] + E_t[u(a_t^*, t)] - c$, which is clearly maximized at $c = 0$. Since the value of information is always weakly positive with quasilinear payoffs, the overall maximum is at $c = 0$ (if the value of information is zero, $c \geq \tilde{c}$ also gives a maximum).

In the case where the principal's payoff is non-increasing, the agent's ex-ante IR constraint is always binding, so that aggregate welfare is equal to the principal's payoff. In this case, welfare must be maximized at $c = 0$ and minimized at high levels of cost $c \geq c^*$:

Corollary 1 *If the principal's payoff is U-shaped in cost then aggregate welfare is minimized at $c = c^*$ and maximized at $c = 0$. If the principal's payoff is non-increasing in cost then aggregate welfare is minimized at $c \geq c^*$ and*

[5]With quasilinear payoffs aggregate welfare is straghtforward to define. In the general case it will be necessary to be more careful in the definition.

maximized at $c = 0$.

In some cases, the principal will always leave rents to the agent for any c in the range $(c^*, \tilde{c})$. The following diagram illustrates the principal's payoff and total welfare when the environment is a unit-valuation monopoly with zero marginal costs. That is $A = [0, 1]$, where $a \in A$ denotes the probability of receiving the good, and payoffs are:

$$U(a, t) = t \cdot a$$

$$V(a, t) = 0.$$

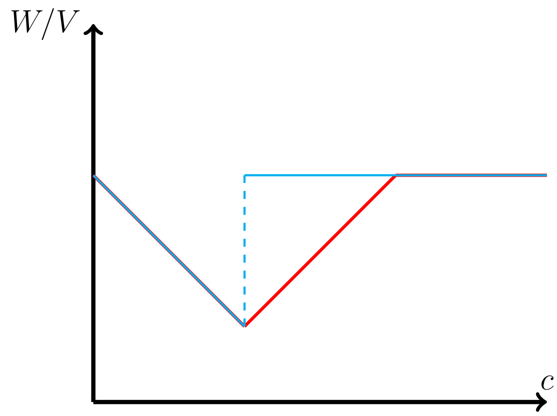

Figure 2: Payoffs in the unit-valuation monopoly model

The blue line represents total welfare: $W = E_t[u(a(t), p(t), t) + u(a(t), p(t), t)] - c\mathbf{1}\{d = Y\}$. The red line represents the expected payoff to the principal. The gap between the blue and red lines therefore represents the agent's payoff. With these payoffs we can show that the principal never wants to distort the action away from first best ($a = 1$) for any c. The value of information in this problem is zero, so that the principal's payoff is U-shaped in cost. When $c \in (c^*, \tilde{c})$ the principal must lower the price to prevent learning, leaving increasing rents to the agent as the cost falls, until the cost reaches c^* at which point the principal prefers to have the agent learn and offer different prices for different realizations of the cost. The principal prefers to lower the price and leave rents to the agent rather than distort the action to prevent learning. Preventing learning through altering the price rather than the allocation is cheaper for the principal in this environment. A key feature of this specific environment that leads to the 'no distortion' result is that the principal is indifferent over the choice of action, a. An analysis of this specific case is contained in the appendix, where I show that the result is qualitatively the same for any signal structure.

2.3.2 General payoffs

The quasilinear case allows for sharper results but excludes some key applications, such as insurance markets. We now return to the general case (where payoffs need not be quasilinear). By similar reasoning to the quasilinear case, the principal's payoff is either U-shaped or monotone non-increasing in cost. This again follows from Lemmas 2 and 3.

To state the results in this section, I introduce some additional notation to describe how the optimal mechanism varies as a function of the cost of evidence. Let $(g(t)(c))_{t \in T}$ denote the optimal vector of allocations offered when evidence is presented on-path and the cost of evidence is c. Let $g(N)(c)$ denote the allocation offered when the agent does not present evidence and the cost is c.

Similarly to the quasilinear case, if evidence is acquired on path, or if evidence is not acquired on path and the cost of evidence is very high, the principal can hold the agent to his reservation utility. For $a \in A$, let $\hat{p}(\alpha)$ solve $E_t[u(\alpha, \hat{p}(\alpha), t)] = E_t[u(0, t)]$. That is, $\hat{p}(\alpha)$ holds the agent to his reservation utility given the lottery over actions α. This price plays a similar role to the role played by $\hat{p}(a)$ in the quasilinear case. However, in the general case, the optimal choice of action may be a lottery, $\alpha \in \Delta(A)$. Let α^* be the ex-ante optimal action. That is:

$$\alpha^* \in argmax_{\alpha \in \Delta(A)} E_t[v(\alpha^*, \hat{p}(\alpha^*), t)].$$

Let $\tilde{c}$ be the threshold cost below which MH is violated by the optimal ex-ante contract, $(\alpha^*, \hat{p}(\alpha^*))$, subject to IR_N.

In general, I define the value of information to the difference between the principal's optimal payoff when the type is common knowledge and the principal's optimal payoff when the type in unknown to both the principal and agent (in the quasilinear case this was simply the sum of the principal's payoff and the agent's payoff):

$$VI = max_{(g(t))_{t \in T}: u(g(t),t) \geq u(0,t) \forall t \in T} E_t[v(g(t), t)]$$
$$- max_{g(N): E_t[u(g(N),t)] \geq E_t[u(0,t)]} E_t[v(g(N), 0)].$$

Since $v(g(t), t)$ is a function of both price and action, these expressions are the generalizations of $E_t[V(a_t, t)+U(a_t, t)]$ and $E_t[V(a, t)+U(a, t)]$, respectively, given that the agent's payoff is held to the outside option. Note that in both cases the equilibrium payoff to the agent is zero, justifying the focus on the principal's payoff. We have the following results:

Proposition 3 *The principal's payoff is non-increasing in c if and only if $E_t[v(g(t)(\tilde{c}), t)] \geq E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$, otherwise it is U-shaped. In particular, if the value of information is non-positive and $u(\alpha^*, \hat{p}(\alpha^*), t) < u(0, t)$ for some $t \in T$ then the principal's payoff is U-shaped in the cost of evidence.*

A proof of Proposition 3 is contained in the appendix and follows similar lines to the proof of Proposition 1. As in the quasilinear case, the agent's payoff is (weakly) highest for intermediate costs of evidence:

Proposition 4 *The agents ex-ante payoff is greater than zero only if c is between c^* and $\tilde{c}$.*

Exactly, as in the separable case, if $c \leq c^*$ then evidence is acquired on path and IR_Y is binding. If $c > \tilde{c}$ then no evidence is acquired on-path, the optimal contract is $(\alpha^*, \hat{p}(\alpha^*))$ and IR_N is binding. The following diagrams illustrate the two cases:

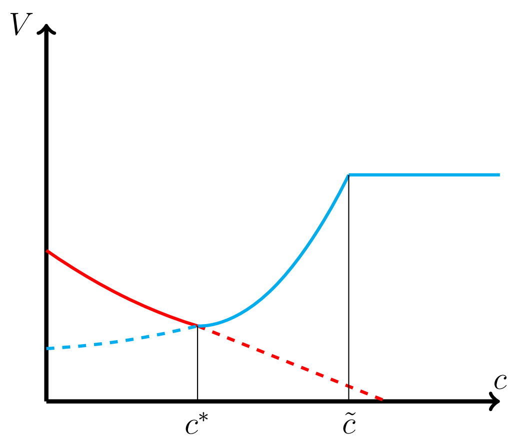

Figure 3: Principal's payoffs: U-shaped case

The red line indicates the optimal payoff from mechanisms that induce evidence acquisition. The blue line indicates the optimal payoff from mechanisms which do not induce evidence acquisition. As we have seen, the former is decreasing in the cost of evidence. The overall optimal payoff is given by the upper envelope of the two lines, indicated by the solid segments.

The flat section at the right on the blue line indicates the payoff when $c > \tilde{c}$, so that $(\alpha^*, \hat{p}(\alpha^*))$ is feasible. In the depicted example, the value of information is negative, so that the red and blue lines cross (at c^*) below the point, $\tilde{c}$ at which the blue line becomes flat. The optimal payoff is therefore U-shaped in the cost of evidence.

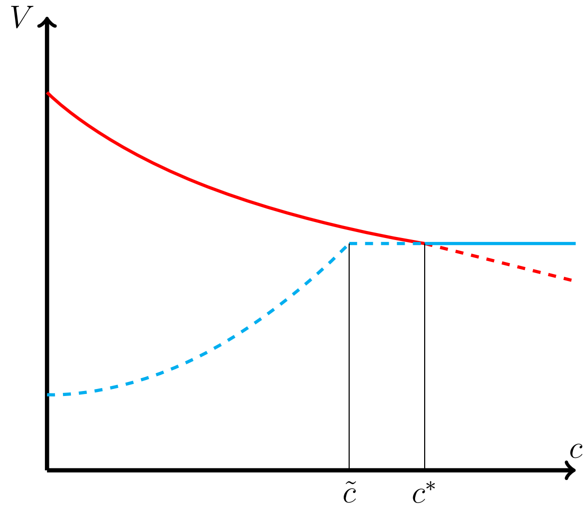

Figure 4: Principal's payoffs: nonincreasing case

In the second case $c^* > \tilde{c}$, so that the principal selects a mechanism that does not induce evidence acquisition only if c is sufficiently high that $(\alpha^*, \hat{p}(\alpha^*))$ is feasible. In this case the principal's payoff is decreasing in the cost of evidence up to c^*, above which it is constant. Again, in the diagram, the blue line represents the optimal payoff from a mechanism which does not induce evidence acquisition and the red line represents the optimal payoff from a mechanism which does induce evidence acquisition. The solid portions indicate the overall optimal payoff: that is, the upper envelope of the two lines.

As in the quasilinear case, aggregate welfare is minimized at c^* in the U-shaped case, and minimized for $c \geq c^*$ in the nonincreasing case: at these values the prinicpal's payoff is lowest and the agent's payoff is held to the outside option. This is true for any definition of aggregate welfare that respects the Pareto order: that is, any measure by which aggregate welfare is lower if both players' payoffs are lower.

The cost for which aggregate welfare is maximized is sensitive to the measure of aggregate welfare used: for example, methods using Pareto weights will give different answers depending on the weights, and methods using a variant of the compensating variation will give different answers depending on how the compensating variation is defined.

The answer to the question "who benefits from the availability of evidence" depends on which of the two cases holds. In the U-shaped case, which occurs when the value of evidence is not too high, the agent may benefit, and the principal may be made worse off by the availability of evidence, if the cost of evidence is in some intermediate range. In the nonincreasing case the principal, and only the principal, benefits from the availability of evidence.

3 Application: Insurance Markets

In this section I develop an application to insurance markets in greater detail. In particular, I ask under what parameter values, in a monopoly insurance model (a la Stiglitz, 1977) does allowing an insurer to base the terms of an insurance contract on disclosure of evidence either (i) make the consumer worse off compared to, or (ii) Pareto dominate the equilibrium when the insurer cannot do so. We will also see that allowing insurers to base the terms of their contracts on disclosure of evidence may either increase or decrease a measure of aggregate welfare.

This comparison sheds light on the policy debate over the regulation of genetic information in insurance contracting. The results show how the validity of the argument that allowing insurers to contract on the results of a genetic test will make consumers worse off depends on the underlying parameters of the market.

The model is similar to the benchmark Stiglitz model with two types, except that the consumer is initially uninformed and can, at a fixed cost, observe a signal that reveals his risk type. If the insurer is allowed to contract on evidence then the signal acts as evidence, as in the general model. Since the value of information in an insurance market is negative (as noted by Hirshleifer, 1971), the insurer's profit is U-shaped in the cost of evidence.

If the insurer is not allowed to contract on evidence then the environment is modeled as a mechanism design problem with (covert) endogenous information acquisition. The model in this case is identical to the model in Lagerlof and Shottmuller (2016a), and will act as a benchmark from which to compare the welfare effects of allowing contracting on evidence.

I find that the consumer is worse off when evidence can be contracted upon if the cost is low and the probability of observing a high signal is low, or if the cost is in an intermediate range. In both cases, the cost is sufficiently low that the insurer will induce evidence acquisition and fully extract ex-ante rents when evidence can be contracted upon. However, if evidence cannot be contracted upon, and parameters are in these ranges, the insurer leaves ex-ante rents to the agent. In the first case, she does so to give the agent incentives to learn; in the second case, to give incentives not to learn.

For some other ranges of parameter values, however, the consumer receives no ex-ante rents even in the case when evidence cannot be contracted upon, while the insurer is strictly better off with the ability to contract on evidence. In these ranges, therefore, allowing the insurer to contract on evidence leads to a Pareto improvement.

The results suggest that whether or not consumers are harmed by allowing insurers to base the terms of contracts on the disclosure of a test result depends on the empirical features of a particular market.

Unlike in other theoretical papers considering this issue, I model a monopolistic, rather than competitive, market. While real life insurance markets often feature some degree of competition, evidence suggests that many insurance markets feature significant competition. Monopoly, therefore, may in some cases be a more realistic starting point for analysis than perfect competition.

3.1 Model

A risk-neutral monopoly insurer faces a risk-averse consumer with an unknown probability of a loss of size L. The consumer has initial wealth $W > L$. The probability of a loss is ϕ_H with probability λ and ϕ_L with probability $1 - \lambda$, where $\phi_L < \phi_H$. We define the ex-ante probability of a loss as $\tilde{\phi} \equiv \lambda\phi_H + (1 - \lambda)\phi_L$.

At a cost, c, the agent can learn his risk type, $t \in \{L, H\}$. I compare two regulatory environments. In the first, the consumer can disclose what he has learned to the insurer and the insurer is free to offer different contracts contingent on disclosure of different risk types (in the case of genetic testing this would correspond to offering different contracts for different test results). This case is modeled as a special case of the general model.

In the second regulatory environment, the insurer must offer the same (menu of) contract(s) to the consumer regardless of any evidence that might be disclosed. For example, there is a law such as the GINA that prohibits contracting on the results of a test. This case is modeled as a contracting problem with covert, non-verifiable information acquisition.[6] This model has been analyzed by Lagerlof and Schottmuller (2016a) and will act as a benchmark

[6] The consumer may be free to disclose test results to the insurer. However, if the insurer cannot contract upon the results, there is no incentive for the consumer to make a disclosure or for the insurer to ask for a disclosure.

against which to compare the welfare effects of allowing contracting on evidence.

As in the general model, contracts can be described by a pair, (a, p), where a is the level of coverage and p is the premium. Ex-post payoffs conditional on risk type $t \in \{L, H\}$ are given by:

$$u(a, p, t) = \phi_t U(W - (1-a)L - p) + (1 - \phi_t)U(W - p)$$

for the consumer, and

$$v(a, p, t) = p - \phi_t a L$$

for the insurer. The function $U(\cdot)$ satisfies $U' > 0$ and $U'' < 0$ — that is, the consumer prefers higher consumption levels and is risk averse. It will be without loss of generality to require that $a \in [0, 1]$.

We denote the ex-ante expected utility from a given contract, (a, p) by $u^\lambda(a, p) \equiv \lambda u(a, p, H) + (1 - \lambda)u(a, p, L)$.

We will denote the value of the outside option for the low, high and ex-ante types by $\tilde{u}$, $\tilde{u}^H$ and $\tilde{u}^\lambda$, respectively. That is:

$$\begin{aligned}\tilde{u} &= \lambda u(0, 0, H) + (1 - \lambda)u(0, 0, L),\\ \tilde{u}^H &= u(0, 0, H),\\ \tilde{u}^L &= u(0, 0, L).\end{aligned}$$

As we have seen in the general model, when the insurer is allowed to contract on evidence there is a threshold cost, $c^* > 0$ below which the insurer prefers to induce evidence acquisition. Since the value of information in insurance markets is negative (cf. Hirshleifer, 1971), it follows from Proposition 3 that the payoff to the insurer is U-shaped in the cost of evidence. Moreover, there is a range of costs $(c^*, \tilde{c})$ for which the agent may obtain positive rents. In fact, in this model the consumer will obtain positive rents for a strict subset of $(c^*, \tilde{c})$, while for all $c \in (c^*, \tilde{c})$ the optimal contract will involve partial insurance and this be distorted away from first best (full insurance).

If evidence cannot be contracted upon, there will also be a threshold cost, c', below which the principal prefers to induce evidence acquisition. This is because if evidence is not acquired then the contract must satisfy the same moral hazard constraint, so that again the insurer's profit from the optimal contract in which evidence acquisition is discouraged is increasing in cost.

If evidence is acquired on path, then the insurer will offer a menu that separates types (otherwise it would be better not to induce evidence acquisition). However, in this case, since the principal cannot demand evidence of the type, the menu must respect incentive compatibility constraints that insure the consumer, having observed a high risk rest result, does not prefer the contract intended for a consumer with a low risk test result. Moreover, the contract must, in expectation, make it worthwhile for the consumer to acquire evidence. It follows that the payoff from the optimal separating menu is decreasing in the cost of evidence. Similarly to the case where evidence can be contracted upon, it follows that a threshold cost, c' exists, and it will also be clear that $c' < c^*$.

Therefore, the insurer is always weakly worse off when she cannot contract on evidence compared to the case in which she can. Moreover, if the insurer would prefer to induce evidence acquisition even when she cannot contract on evidence, then she will also prefer to induce evidence acquisition when she can contract on evidence:

Proposition 5 *There exists a threshold cost, c^* below which the principal is strictly better of when she can contract on evidence. If the cost is above c^* then the optimal contract, and the principal's payoff, are the same whether she can contract on evidence or not.*

The main results of this section characterize parameter values for which the consumer is strictly better off when evidence cannot be contracted upon, and parameter values for which the allowing the principal to contract upon evidence induces a Pareto improvement — the consumer is no worse off, while the consumer is strictly better off.

Lagerlof and Schottmuller's analysis shows that the consumer receives positive rents if (i) evidence is acquired on-path and the probability of the high type is high or (ii) evidence is not acquired on-path and the cost is below a threshold, c''. If the cost is greater than c^*, so that evidence is not acquired on-path even when it can be contracted upon, the equilibrium contract is the same under both regulatory regimes. I will show that, when evidence can be contracted upon, the consumer receives positive rents only if $c^* < c < c''$. Therefore, we can show that:

Proposition 6 *The agent is better of when evidence cannot be contracted upon if and only if the cost of evidence is low and the ex-ante probability of a high type is low or if the cost is in a (possibly empty) intermediate range.*

Formally, there exist threshold costs, $c^* > 0$ and $c'' \geq c' > 0$, and a threshold probability, $\lambda^* > 0$, such that the consumer is better off when evidence cannot be contracted upon if and only if $c < c'$ and $\lambda < \lambda^*$, or $c \in [c', c''] \cap [0, c^*]$.

These propositions lead to the following interesting corollary:

Corollary 2 *There exist ranges of parameters for which the optimal contract when evidence can be contracted upon Pareto improves on the optimal contract when evidence cannot be contracted upon. In particular, when* $c < \tilde{c}$ *and* $\lambda \geq \lambda^*$ *or if* $c \in [c^*, \tilde{c}] \cap (c', \infty)$ *then the principal is strictly better off when evidence can be contracted upon, while the agent is no worse off.*

To prove the proposition I will first review the benchmark case, in which evidence cannot be contracted upon, and then characterize the optimal mechanism when evidence can be contracted upon. The results then follow by comparing the payoffs implied by the optimal mechanisms in the two cases.

The characterization of the optimal mechanism when evidence can be contracted upon is also of independent interest as it contrasts with a key finding in the literature on mechanism design with endogenous (non verifiable) information: that mechanisms with endogenous information are especially sensitive to the agent's type. In contrast, I will show that the contracts offered after endogenous evidence is presented is less sensitive to the agent's type than the contracts that would be offered if evidence was exogenously given.

3.2 Optimal mechanism when evidence cannot be contracted upon

When evidence cannot be contracted upon, the insurer faces a choice between a mechanism which does not induce evidence acquisition (which, as in the case with contractible evidence, must satisfy a moral hazard constraint ensuring that the consumer does not prefer to learn his type and purchase only if the type is high), and a mechanism which does induce evidence acquisition.

The difference is that (effectively) types cannot be credibly disclosed once they have been learned. Therefore, if evidence is acquired on path the principal will have to offer a menu of contracts among which the agent self-selects. In fact, an additional incentive constraint arises that is not present in the case of exogenous private information: the agent could always *not* learn his type and purchase one of the contracts independently of the ex-post type. [7] The need to prevent the agent from doing so further constrains the set of feasible menus.

As in the case with contractible evidence, the optimal mechanism can be derived by comparing the optimal mechanism in which evidence is acquired on-path with the optimal mechanism that discourages the agent from learning. We first consider the optimal mechanism that discourages learning.

The problem is a specific case of the problem analyzed in the general model. In particular, the principal solves:

$$\begin{aligned} \text{maximize} \quad & p - \tilde{\phi} a \\ \text{subject to} \quad & u^{\lambda}(a,p) \geq \tilde{u} \ \ (IR_N), \\ & (1-\lambda)(\tilde{u}^{L} - u(a,p,L)) \leq c \ \ (MH). \end{aligned}$$

The particular form of the MH constraint is justified by the fact that, as long as IR_N is satisfied, the high type will continue to purchase even after learning the risk. This is because insurance is valued more conditional on high risk than it is ex-ante. Notice that this constraint trades off the value of the contract and the value of the outside option according to the *low type's* preferences. This is intuitive: the difference between the gross payoff if the consumer always purchases insurance and the gross payoff if the consumer only purchases when the risk is high is the difference between the value of the contract and the outside option when the risk is low.

[7] In the case where evidence can be contracted upon this is not possible because choosing a contract requires disclosing evidence.

Lagerlof and Schottmuller (2016) show that the best contract for the insurer for which the consumer does not acquire evidence is the following:

Lemma 5 *There exist threshold costs $c'' \geq 0$ and $\tilde{c} > c''$ such that the optimal contract with no on-path learning is given by:*

$$(a,p) = \begin{cases} (1, W - U^{-1}(\tilde{u})) & c \geq \tilde{c} \\ (a^*, p^*) : u^{\lambda}(a,p) - \tilde{u} = (1-\lambda)(\tilde{u}^L - u(a,p,L)) - c = 0 & c \in (c'', \tilde{c}) \\ argmax_{(a,p)} p - \tilde{\phi} a \;\; s.t. \; (MH) & c \leq c''. \end{cases}$$

In the first case, $c \leq c'''$, the cost is sufficiently high that the first best — full insurance at the ex-ante certainty equivalent — is feasible. In the second case, the principal selects a contract at which both IR_N and MH are binding. The contract involves partial insurance. In the third case, when $c \leq c'''$, the first best is not feasible, but the solution to the relaxed problem in which only MH is taken into account is feasible, and therefore optimal. In this third case the agent receives positive rents.

The optimality of this contract can be explained intuitively with a diagrammatic sketch:

In the following diagrams contracts are represented in the space of ex-post consumption pairs, (x_1, x_2), where x_1 is consumption in the loss state and x_2 is consumption in the state where no loss occurs. The agent prefers contracts further from the origin, the principal prefers contracts further from the origin.

The following diagram depicts the optimal contract when $c > \tilde{c}$:

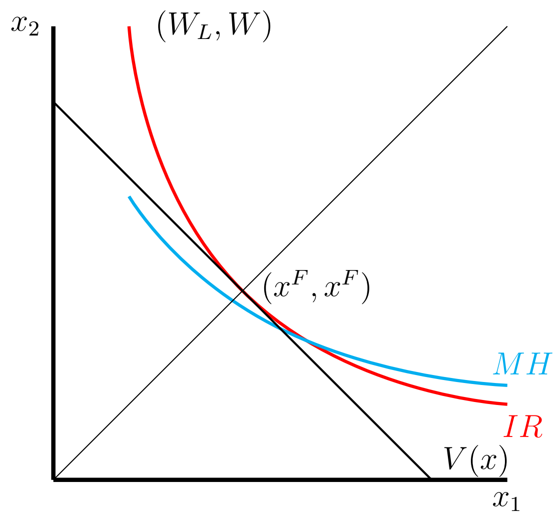

Figure 5: Optimal contract when $c \geq \tilde{c}$

The contract (x^F, x^F) is the first best contract, subject to IR_N. If the cost is sufficiently high, then MH is satisfied at this contract, therefore it is optimal. The threshold cost $\tilde{c}$ above which this contract is optimal is strictly above zero, since this contract makes the consumer indifferent to the outside option ex-ante. It follows that if the type could be learned at zero cost, the consumer would prefer not to purchase this contract conditional on the low risk type.

If the first best is not feasible, a second best contract would maximize the insurer's profit subject to MH only. Clearly, as the solution to a relaxed problem, this contract is optimal if feasible. It can be shown that for sufficiently low cost (possibly, but not necessarily, only $c = 0$) this second best contract is in fact feasible. This case is illustrated in the following diagram:

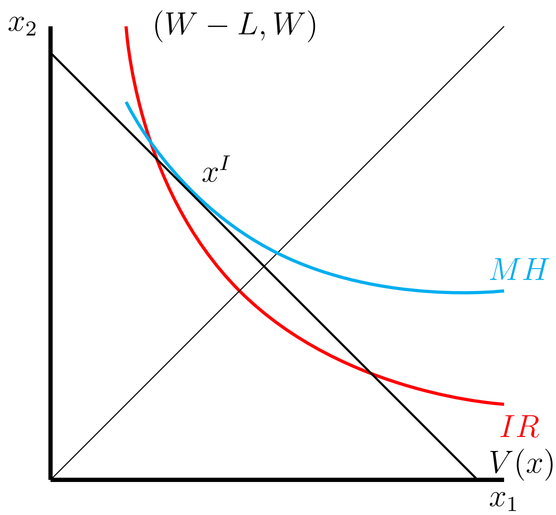

Figure 6: Optimal contract when $c < c''$

If $c \in (c'', \tilde{c})$ then neither (x^F, x^F) nor (x_1^I, x_2^I) is feasible. In this case the principal would like to choose the closest feasible point to full insurance along the set of points at which the IR constraint is binding. This follows because, since the consumer is risk averse while the insurer is risk neutral, increasing the level of insurance while leaving the consumer's ex-ante payoff fixed increases the insurer's profit. That is, for intermediate costs, the optimal contract, (a^*, p^*) satisfies:

$$(1-\lambda)[\tilde{u}^L - u^L(a^*, p^*)] - c = \tilde{u} - u^\lambda(a^*, p^*) = 0$$

so that both constraints are binding. The solution in this case is illustrated in the following diagram:

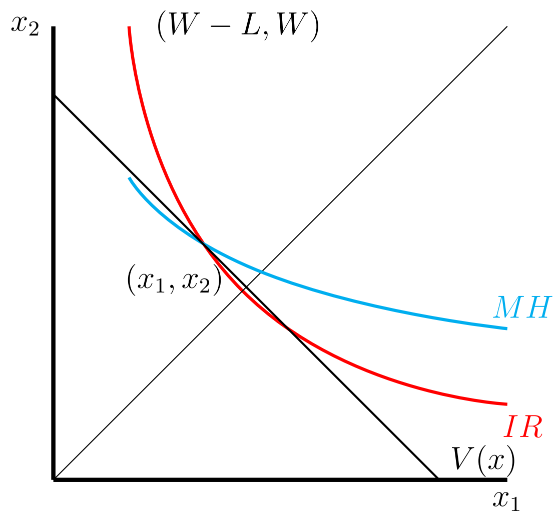

Figure 7: Optimal contract when $c'' < c < \tilde{c}$

Note that, for all subcases except for $c < c''$, the IR constraint is binding. We therefore have the following corollary:

Corollary 3 *If evidence is not acquired on-path, the agent receives positive ex-ante rents if and only if $c < c''$.*

We can also show that the solution implies partial insurance unless the first-best is feasible:

Corollary 4 *If evidence is not acquired on-path and $c < \tilde{c}$ then $a < 1$ in an optimal contract.*

This is clear from the diagrammatic sketch.

We next consider the case in which evidence is acquired on-path but the insurer cannot contract upon it. In this case, the insurer will offer a menu of contracts. Note that the insurer will never offer a pooling menu, since in this case there would be no incentive for the agent to acquire evidence.

The optimal menu, $(a(t), p(t))_{t=H,L}$, solves the following problem:

$$\begin{aligned}
&\text{maximize } \ \lambda(p(H) - \phi_H La(H)) + (1-\lambda)(p(L) - \phi_L La(L)) \\
&\text{subject to:} \\
\lambda u(a(H), p(H), H) + (1-\lambda)u(a(L), p(L), L) &\geq \tilde{u} + c && (IR_Y), \\
u(a(H), p(H), H) &\geq u(a(L), p(L), H) && (IC_H), \\
u(a(L), p(L), L) &\geq u(a(H), p(H), L) && (IC_L), \\
u(a(H), p(H), H) &\geq \tilde{u}^H && (IR_L), \\
u(a(L), p(L), L) &\geq \tilde{u}^L && (IR_L), \\
\lambda u(a(H), p(H), H) + (1-\lambda)u(a(L), p(L), L) - c &\geq \\
\lambda u(a(L), p(L), H) + (1-\lambda)u^L(a(L), p(L), L) & && (MH_L), \\
\lambda u(a(H), p(H), H) + (1-\lambda)u(a(H), p(H), H) - c &\geq \\
\lambda u(a(H), p(H), H) + (1-\lambda)u(a(H), p(H), L) & && (MH_H).
\end{aligned}$$

The problem is different to the principal's problem with on-path evidence acquistion in the general environment, because in this case evidence cannot be contracted upon. As in the standard benchmark (Stiglitz, 1977), the menu of contracts must satisfy IC constraints ensuring that each ex-post type of the agent self selects into the appropriate contract. However, additional constraints are imposed (MH_L and MH_H) because the agent may choose not to learn his type and simply select and element from the menu ex-ante. Doing so would not be detectable by the principal since, as in the whole of the present paper, the evidence acquisition decision is not observable.

These additional constraints are in fact stricter than the standard incentive compatibility constraints. For example, a simple manipulation shows that MH_L is equivalent to:

$$u(a(H), p(H), H) \geq u(a(L), p(L), L) + \frac{c}{\lambda}.$$

So that IC_H is implied by MH_L. Similarly, IC_L is implied by MH_H. Nevertheless, Lagerlof and Schottmuller show that the optimal mechanism is qualitatively similar to the optimal mechanism in the standard case analyzed by Stiglitz (1977): if the probability of the high type is high, then the low type is excluded and the high type receives a full insurance contract which leaves IR_Y binding (note that this implies IR_H is slack. If the probability of the high type is low, then the low type receives a partial insurance contract with IR_L binding, the high type receives a full insurance contract, and MH_H is binding. In the latter case, the agent receives positive rents ex-ante: that is, IR_Y is not binding.

For our results the important implication is the following:

Lemma 6 *If evidence cannot be contracted upon and is aquired on-path, then there exists a treshold probability, $\lambda^* > 0$, such that IR_Y is slack given the optimal contract if and only if $\lambda < \lambda^*$.*

As we have noted earlier, evidence will be acquired on-path if and only if the cost of evidence acquisition is below some threshold, $c' > 0$. Therefore, the overall optimal mechanism is the mechanism described in Lemma 5 if $c \geq c'$ and the mechanism described in Lemma 6 if $c < c'$. Lagerlof and Schottmuller show that c' may be larger or smaller than c'' depending on the parameter values. The payoff implications, which will form a benchmark against which we compare the payoffs from the optimal mechanism when evidence can be contracted upon are summarized by the following proposition:

Proposition 7 *There exist threshold costs $c' > 0$ and $c'' \geq 0$ and a threshold probability $\lambda^* > 0$ such that the agent receives positive rents from the optimal mechanism when evidence cannot be contracted upon if and only if $c \geq c'$ and $\lambda \geq \lambda^*$ or $c \in (c', c'')$.*

3.3 Optimal mechanism when evidence can be contracted upon

In case 3 evidence is acquired on-path and the insurer offers a different contract after each realization is disclosed (and offers the outside option if no evidence is disclosed). Ideally the principal would like to offer a pooling contract that fully extracts the full surplus ex-ante, net of evidence acquisition costs. Since the value of evidence ex-ante is greater than the value of evidence ex-post (cf. Hirshleifer, 1971), this is best for the principal if we only take ex-ante IR into account. However, since the agent, having acquired evidence, knows the risk type, this will generally not be feasible. In fact, if the low type would accept this contract then it would also be feasible without evidence, in which case the principal could generally benefit from the reduction in cost. It will follow that whenever evidence is used on-path, a different contract will be offered to the low and high risk types. Nevertheless, both contracts will feature full ex-post insurance:

Lemma 7 *Suppose that evidence is acquired on-path and can be contracted upon. Then* $a(L) = a(H) = 1$.

Suppose not. For example, suppose that $a(L) < 1$. Then a new contract, $(a'(L), p'(L))$ with $a'(L) = 1$ and $u(a'(L), p'(L), L) = u(a(L), p(L), L)$ will improve the insurer's profit (since the insurer is risk neutral while the consumer values insurance) and leave all relevant constraints satisfied (since the IR contraints – the only relevant constraints in this case – depend only on the interim payoffs $u^L(\cdot)$ and $a^L(\cdot)$.

As we saw in the general case, we must also have the ex-ante IR constraint binding. If it is not, then either IR_L or IR_H must be slack, so that the insurer's profit can be improved by increasing either $p(L)$ or $p(H)$.

Moreover, we always have $p(L) < p(H)$. Suppose not, then $u(a(H), p(H), H) \geq u(a(L), p(L), L) \geq \tilde{u}^L > \tilde{u}^H$ so that IR_H is not binding. Then consider decreasing $p(L)$ and increasing $p(H)$ slightly keeping $EAIR$ binding. This is feasible and must improve the principal's profit as it provides more ex-ante insurance (by risk aversion of the agent).

In addition to $EAIR$ binding, at least one of the ex-post IR constraints must bind at an optimum:

Lemma 8 *Either IR_H or IR_L is binding.*

Suppose to the contrary that neither is binding and $p(L) < p(H)$. Then, we can increase $p(L)$ and decrease $p(H)$ keeping $EAIR$ binding. As we have seen in previous cases, this increases profit (since the consumer values insurance ex-ante).

We are left with two possibilities. Either $p(H) = W - U^{-1}(\tilde{u}^H)$ and $p(L) = W - U^{-1}(\tilde{u}^L + \frac{c}{1-\lambda})$ or $p(L) = W - U^{-1}(\tilde{u}^L)$ and $p(H) = W - U^{-1}(\tilde{u}^H + \frac{c}{\lambda})$. We can show that, in fact, the latter always does best. Intuitively, it provides more ex-ante insurance for the same ex-ante consumer surplus, hence higher insurer profits. A proof is contained in the appendix.

To summarize:

Lemma 9 *Suppose that evidence is acquired and presented on-path. Then the equilibrium contract is full ex-post insurance ($a(L) = a(H) = 1$) with:*

$$p(L) = W - U^{-1}(\tilde{u}^L)$$
$$p(H) = W - U^{-1}(\tilde{u}^H + \frac{c}{\lambda}).$$

Note that, with exogenous evidence, the optimal set of contracts would be full insurance ex-post with $p(H) = W - U^{-1}(\tilde{u}^H)$ and $p(L) = W - U^{-1}(\tilde{u}^L)$. With endogenous evidence we have $p(H) - p(L) < W - U^{-1}(\tilde{u}^H) - (W - U^{-1}(\tilde{u}^L))$, so that in this sense the optimal menu of contracts is 'low powered': the contract varies less with endogenous evidence than with exogenous evidence. This contrasts with the findings of the endogenous information literature (such as Cremer et al., 1998) which finds that optimal contracts with non verifiable endogenous information features contracts that are especially sensitive to information.

3.4 Overall solutions and welfare comparisons

Having characterized the optimal mechanism in cases 1-3 we can prove the proposition by characterizing the overall optimal mechanism when the principal can or cannot contract on evidence and comparing the principal's and agent's payoff between the two cases.

We first note that, both when the principal can contract upon evidence and when the principal cannot, there is a threshold cost below which the principal induces evidence acquisition on-path; the threshold is $c' > 0$ when the principal cannot contract upon evidence and $c^* > 0$ when the principal can contract upon evidence. The existence of c^* was established in the general model, the existence of c' is proved by Lagerlof and Schottmuller.

In both cases, the existence of a threshold cost follows because the insurer's profit is decreasing in c if evidence is acquired on-path and increasing in c if is it not.

Moreover, we have that $c' < c^*$ because the profit from a separating pair of contracts when evidence can be contracted upon is always greater than the profit from a separating menu when evidence cannot be contracted upon. That the profit is strictly weakly greater follows because the optimal mechanism with on-path evidence acquisition must respect fewer constraints when evidence can be contracted upon. That the profit is strictly greater follows because the solution when evidence can be contracted upon is unique. Therefore, the set of parameters for which evidence acquisition is better for the principal is strictly larger when evidence can be contracted upon.

When evidence can be contracted upon and is acquired on-path the optimal set of contracts is the set characterized in Lemma 9. When evidence cannot be contracted upon and is acquired on-path, the opitmal menu is the menu described in Lemma 6. Whether or not evidence can be contracted upon, if it is not acquired on-path then the optimal contract is the contract characterized in Lemma 5.

In particular, when $c > c^*$ the contract offered by the principal induces no acquisition, and is the same whether or not evidence can be contracted upon. If $c < c^*$ the pair of contracts offered when evidence can be contracted upon leaves no rent to the agent. Therefore, the agent is strictly better of when evidence cannot be contracted upon if and only if $c < c^*$ and the contract offered when evidence cannot be contracted upon leaves positive rents to the agent. It follows from Proposition 7 that:

Lemma 10 *The agent is strictly better off when evidence cannot be contracted upon if* $c \in [c', c^*)$ *and* $c < c''$*, or if* $c < c'$ *and* $\lambda < \lambda^*$ *. Otherwise, the agent obtains zero rents whether or not evidence can be contracted upon.*

Meanwhile, the insurer's payoff, it is strictly higher when evidence can be contracted upon as long as evidence is presented on-path (that is, $c < c^*$), otherwise it is the same in both cases. This follows because the contract offered when evidence cannot be contracted upon is different to the contract offered when evidence can be contracted upon. Meanwhile the contract offered when evidence can be, and is, contracted upon is the unique solution to a problem with fewer constraints.

The main propositions of this section follow from these observations.

Aggregate welfare may also be lower when evidence can be contracted upon in the following sense: suppose that $[c', c^*)$ is nonempty, and the cost of evidence acquisition is slightly below c^*. At c^* the payoff to the insurer is the same in both contracting environments. Since the insurer's payoff changes continuously with cost in both contracting environment, at $c = c^* - \epsilon$ (for small ϵ) the difference in the insurer's payoff between the case in which she can contract on evidence and the case in which she cannot is arbitrarily small. On the other hand, the agent's payoff is lower by a discrete amount when evidence can be contracted upon. It follows that if aggregate welfare is measured as a weighted sum of payoffs then, for any positive weight on the agent's payoff, aggregate welfare is lower when evidence can be contracted upon for some cost $c = c^* - \epsilon$.

Intuitively, this occurs because evidence acqusition is costly but socially unproductive. When the insurer can contract upon evidence, she has more incentive to induce evidence acquisition because it may (depending on parameter values) allow her to extract a larger fraction of surplus. However, when evidence is acquired the consumer incurs a socially unproductive acquisition cost and market efficiency may be reduced due to the Hirshleifer effect. Therefore, aggregate welfare can be reduced. For other parameter values (for example those where allowing contracting on evidence leads to a Pareto improvement), the effect on aggregate welfare will be in the opposite direction.

The finding that the consumer is weakly worse off (and sometimes strictly worse off) when evidence can be contracted upon relies on the insurer having market power. Doherty and Thistle (1996) study a similar model in the context of a competitive market. Their findings suggest that the consumer is no worse off, and sometimes better off, when evidence can be contracted upon.

Proposition 3 in their paper shows that, when evidence cannot be contracted upon, either no evidence is acquired or an equilibrium does not exist. If evidence can be contracted upon, in the monopolistic model, the consumer must be better off, ex-ante, than in the best feasible pooling contract, otherwise deviating to a pooling contract would be a profitable deviation for an insurer. Together, these observations imply that the consumer is better off when evidence can be contracted upon. Indeed, for some parameter values (in particular, low costs of evidence acquisition) an equilibrium exists only if evidence can be contracted upon.

A primary concern in the debate over the GINA is that consumers may be made worse off by the ability of insurance companies to contract on evidence. These results suggest that this concern is valid only if market power is a concern in the given market. I have modeled market power using the extreme case of monopoly power, however the intuition may carry over to less extreme forms of market power. An additional concern, which is highlighted above, is that allowing insurers to base the terms of their contracts on the disclosure of evidence may make the market less *efficient* as well as simply making consumers worse off. Again, this is a concern only if market power is present. Given empirical evidence that substantial market power does exist in insurance markets, these findings are relevant to the policy debate.

4 References

Barigozzi, Francesca and Dominique Henriet (2011). "Genetic Information: Comparing Alternative Regulatory Approaches When Prevention Matters." *Journal of Public Economic Theory*, 13, pp. 23-46.

Barzel, Yoram (1977). "Some Fallacies in the Interpretation of Information Costs." *Journal of Law and Economics*, 20, pp. 291-307.

Ben-Porath, Elchannan, Eddie Dekel and Barton L. Lipman (2017). "Mechanisms with Evidence: Commitment and Robustness." Mimeo.

Bergemann, Dirk, Benjamin Brooks and Stephen Morris (2015). "The Limits of Price Discrimination." *The American Economic Review*, 105, pp. 921-957.

Bergemann, Dirk and Juuso Valimaki. "Information Acquisition and Efficient Mechanism Design." *Econometrica*. 70, pp. 1007-1033.

Chiappori, Pierre-Andre, Bruno Julien, Bernard Salanie and Francois Salanie (2006). "Asymmetric Information in Insurance: General Testable Implications." *The RAND Journal of Economics*, 37, pp. 783-798.

Cohen, Alma and Liran Einav (2007). "Estimating Risk Preferences from Deductible Choice." *The American Economic Review*, 97, pp. 745-788.

Craswell, Richard (1988). "Precontractual Investigation as an Optimal Precaution Problem." *The Journal of Legal Studies*, 17, pp. 401-436.

Cremer, Jacques and Fahad Khalil (1992). "Gathering Information Before Signing a Contract." *The American Economic Review*, 82, pp. 566-578.

Cremer, Jacques and Fahad Khalil (1994). "Gathering Information Before the Contract is Offered: The Case with Two States of Nature." *European Economic Review*, 38, pp. 675-682.

Cremer, Jacques, Fahad Khalil and Jean-Claude Rochet (1998a) "Contracts and Productive Information Gathering" *Games and Economic Behavior*, 25, pp. 174-193.

Cremer, Jaques, Fahad Khalil and Jean-Claude Rochet (1998b) "Strategic Information Gathering Before a Contract is Offered." *Journal of Economic Theory*, 81, pp. 163-200.

Dafny, Leemore S. (2010). "Are Health Insurance Markets Competitive?" it The American Economic Review, 100, pp. 1399-1431.

DeMarzo, Peter M., Ilan Kremer and Andrzej Skrzypacz (2015). "Test Design and Minimum Standards." Mimeo.

Demski, Joel S. and David E. M. Sappington (1987). "Delegated Expertise." *Journal of Accounting Reseach*, 25, pp.68-89.

Dang, Tri Vi, Gary Gorton and Bengt Holmstrom (2015). "The Information Sensitivity of a Security." Mimeo.

Doherty, Neil A. and Paul D. Thistle (1996). "Adverse Selection with Endogenous Information in Insurance Markets." *Journal of Public Economics*, 63, pp. 83-102.

Erwin, Cheryl (2008). "Legal Update: Living with the Genetic Information Nondiscrimination Act." *Genetics in Medicine*, 10, pp. 869-873.

Glazer, Jacob and Ariel Rubinstein (2004). "On Optimal Rules of Persuasion." *Econometrica*, 72, pp. 1715-1736.

Glazer, Jacob and Ariel Rubinstein (2006). "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach." *Theoretical Economics*, pp. 395-410.

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

Hirshleifer, Jack (1971). "The Private and Social Value of Information and the Reward to Inventive Activity." *The American Economic Review*, 61, pp. 561-574.

Hoel, Michael, Tor Iversen, Torre Nilssen and Jon Vislie (2006). "Genetic Testing in Competitive Insurance Markets with Repulsion from Chance: A Welfare Analysis." *Journal of Health Economics*, 25, pp. 847-860.

Hoppe, Eva I. and Patrick W. Schmitz (2010). "The Costs and Benefits of Additional Information in Agency Models with Endogenous Information Structures." *Economics Letters*, 107, pp. 58-62.

Hoy, Michael and Mattias Polborn (2000). "The Value of Genetic Information in the Life Insurance Market." *Journal of Public Economics*, 78, pp. 235-252.

Lagerlof, Johan N. M. and Christoph Schottmuller (2016a). "Facilitating Consumer Learning in Insurance Markets – What Are the Welfare Effects."

Mimeo.

Lagerlof, Johan N. M. and Christoph Schottmuller (2016b). “Monopoly Insurance and Endogenous Information.” Mimeo.

Lizzeri, Alessandro (1999). “Information Revelation and Certification Intermediaries.” *RAND Journal of Economics*, 26, pp. 93-113.

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

Persico, Nicola (2000). “Information Acquisition in Auctions.” *Econometrica*, 68, pp. 135-148.

Peter, Richard, Andreas Richter and Paul Thistle (2017). “Engodenous Information, Adverse Selection and Prevention: Implications for Genetic Testing Policy.” *Journal of Health Economics*, 55, pp. 95-107.

Pram, Kym (2017). “Hard Evidence and Welfare in Adverse Selection Environments.” Mimeo.

Roesler, Anne-Katrin and Balazs Szentes (2017). “Buyer-Optimal Learning and Monopoly Pricing.” *The American Economic Review*, 107, pp. 2072-2080.

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

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

Shi, Xianwen (2012). "Optimal Auctions with Information Acquisition." *Games and Economic Behavior*, 74, pp. 666-686.

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

Strohmenger, R. and A. Wambach (2000). "Adverse Selection and Categorical Discrimination in the Health Insurance Markets: The Effects of Genetic Tests." *Journal of Health Economics*, 19, pp. 197-218.

Szalay, Dezso (2009). "Contracts with Endogenous Information." *Games and Economic Behavior*, 65, pp. 586-625.

Terstiege, Stefan (2012). "Endogenous Information and Stochastic Contracts." *Games and Economic Behavior*, 76, pp. 535-547.

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

5 Appendix

5.1 Omitted Proofs from Section 2

Proof of Lemma 1: Firstly suppose that the agent does not randomize over the evidence acquisition decision. We will show that there exists an optimal mechanism out of the set of mechanisms that do not induce evidence acquisition and an optimal mechanism out of the set of mechanisms that do induce evidence acquisition. It follows that there exists an overall optimum given the initial hypothesis of no randomization.

Firstly consider mechanisms that do not induce evidence acquisition. It must be that:

$$E_t[u(g(N),t)] \geq E_t[max\{u(0,t),u(g(t),t),u(g(N),t)] - c$$

and that

$$E_t[u(g(N),t)] \geq u(0,t).$$

Note that the first constraint is always weakly relaxed by letting $g(t) = 0$. Suppose that $g(N)$ is a joint lottery over $A \times P$. Then letting $g'(N) = (marg_A g(N), E_{g(N)}p)$ is an improvement for the principal and relaxes both constraints, under assumption 4. Therefore, without loss of optimality, $p(t)$ is deterministic.

For a given $\alpha \in \Delta(A)$ let $\hat{p}(a)$ satisfy $E_t[u(\alpha,\hat{p}(\alpha),t)] = 0$ and $p_{MH}(\alpha)$ satisfy $E_t[u(\alpha,p_{MH}(\alpha),t)] = E_t[max\{0,u(\alpha,p_{MH}(\alpha),t)\}] - c$. Under assumption 5 $\hat{p}(\alpha)$ exists. To see that, note that p_{MH} is a zero of the expression:

$$E_t[u(\alpha,p,t)] - E_t[max\{0,u(\alpha,p,t)\}] + c$$

Which is positive for arbitrarily low p and negative for arbitrarily high p (by assumption 5), as well as being monotone decreasing and continuous in p (by assumptions 1 and 3).

Let $p^*(\alpha) = min\{\hat{p}(\alpha), p_{MH}(\alpha)\}$. Then $p^*(\alpha)$ is a continuous function and the optimal payoff for the principal out of mechanisms that specify the distribution over actions α and do not induce evidence acquisition is:

$$V^*(\alpha) = E_t[v(\alpha,p^*(\alpha)),t)].$$

Therefore an optimal mechanism is given by $g(N) = (\alpha^*, p^*(\alpha^*))$ where α^*

solves:

$$argmax_{\alpha \in \Delta(A)} V^*(\alpha).$$

Let $\Delta(A)$ be endowed with the weak* topology. Then α^* is the maximizer of a continuous function on a compact set, therefore such an α^* exists.

Now consider mechanisms that do induce evidence acquisition. By similar reasoning to the above, without loss of optimality we can set $g(N)$ equal to the outside option.

The constraints become:

$$E_t[u(g(t), t)] - c \geq 0.$$

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

Similarly to the previous case, without loss of optimality, $p(t)$ is deterministic, so we can write $g(t) = (\alpha(t), p(t))$. Let $\bar{p}(t)$ satisfy $\bar{p}(t) = sup \cup_{\alpha \in \Delta(A)} \{p : u(\alpha, p, t) \geq 0\}$. By assumptions 1 and 5, $\bar{p}(t) < \infty$ for all $t \in T$.

Let $\underline{p}(t) = inf \cup_{\alpha \in \Delta(A)} \{p : u(\alpha, p, t) \geq c/\pi(t)\}$. The second constraint requires that $p(t) \leq \bar{p}(t)$ and without loss of optimality $p(t) \leq \underline{p}(t)$ (since then the entire cost of evidence is reimbursed conditional on type t, given that the ex-post IR constraint is satisfied for all t). Therefore we can take the vector $(g(t))_{t \in T}$ to be contained in a bounded subset of $\Delta(A \times P)$.

Next we show that the feasible set is closed. Suppose that a sequence of vectors $(g(t))^k_{t \in T}$ satisfies the constraints for each $k \in \mathbb{N}$ and $g(t)^k \to g(t)$ for each $t \in T$ (i.e. the sequence of mechanisms converges in the product topology). Then by continuity of u, the limit mechanism satisfies the constraints.

Therefore, the principal's problem consists of maximizing a continuous function over a compact set. A solution therefore exists.

An overall solution exists as the maximum of the two solutions.

Next I show that, without loss of optimality, the agent does not randomize over the evidence acquisition decision. Suppose that he does. Then the agent must be indifferent between acquiring or not acquiring evidence, that is:

$$E_t[u(g(t),t)] - c = E_t[u(g(N),t)].$$

Suppose that $E_t[v(g(t),t)] > E_t[v(g(N),t)]$. Then let $g'(t) = (1-\epsilon)g(t) + \epsilon g(N)$ for small ϵ. Then in the mechanism $(g'(t))_{t\in T}, g(N))$, the agent strictly prefers to acquire evidence and so will do so with probability 1. Moreover, it is easy to see that all constraints continue to be satisfied. Since $E_t[v(g(t),t)] > E_t[v(g(N),t)]$ this is better for the principal as long as ϵ is small.

Conversely suppose that $E_t[v(g(t),t)] \leq E_t[v(g(N),t)]$. Then let $g'(t) = (1-\epsilon)g(t) + \epsilon 0$, where 0 denotes the outside option. Since we must have $E_t[u(g(t),t)] - c \geq E_t[u(0,t)] \Rightarrow E_t[u(g(t),t)] > E_t[u(0,t)]$ (unless $c = 0$), in the mechanism $(g'(t))_{t\in T}, g(N))$, the agent will choose not to acquire evidence with probability 1. Again, by our hypothesis this is better for the principal and all constraints continue to be satisfied.

Finally, consider the remaining case: $c = 0$ and $E_t[v(g(t),t)] \leq E_t[v(g(N),t)]$. In this case, the MH constraint imples that $u(g(N),t) \geq u(0,t)$ for all $t \in T$. Modify $g(t)$ to $g'(t) = g(N)$ for all $t \in T$. Then the principal's payoff is at least as high in the new mechanism and it is easy to see that all constraints are satisfied. The agent is indifferent to acquiring or not acquiring evidence, however since the outcome in the new mechanism is the same in either case, the outcome is unaffected if we assume that the agent acquires evidence for sure.

Proof of Lemma 2: For $\alpha \in \Delta(A)$, let $\hat{p}(\alpha)$ satisfy $E_t u(\alpha, \hat{p}(\alpha), t) = E_t u(0, t)$. Let α^* maximize $E_t[v(\alpha, \hat{p}(\alpha), t)]$. I claim that $E_t[(v(g(N), t)] \leq E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$ for all feasible mechanisms, g, in which the agent does not acquire evidence on-path.

Let g be a feasible mechanism in which the agent does not acquire evidence on path. Without loss of optimality, $g(t) = 0$ for all $t \in T$. Let $g'(N) = (marg_A g(N), E_g(p))$. By Assumption 5, $E_t[v(g'(N), t)] \geq E_t(v(g(N), t)]$. Moreover, by construction, $E_t[v(marg_A g(N), \hat{p}(marg_A g(N)), t)] \geq E_t[v(g'(N), t)]$. By definition, $E_t[v(\alpha^*, \hat{p}(\alpha^*), t)] \geq E_t[v(marg_A g(N), \hat{p}(marg_A g(N)), t)] \geq E_t[v(g(N), t)]$, which proves the claim.

It follows that, out of all mechanisms that do not induce evidence acquisition, $(\alpha^*, \hat{p}(\alpha^*))$ is optimal if feasible. The IR_N constraint is satisfied by construction, so feasibility requires that:

$$E_t[u(\alpha^*, \hat{p}(\alpha^*), t)] = E_t[u(0, t)] \geq E_t[max\{u(0, t), u(\alpha^*, \hat{p}(\alpha^*), t)\}] - c.$$

Clearly this is satisfied if and only if c is weakly greater than some threshold $\tilde{c} \geq 0$.

Let $g^Y(c)$ be the optimal mechanism in which evidence acquired on-path and $g^N(c)$ be the optimal mechanism in which evidence is not acquired on-path. That is, $g^Y(c)$ maximizes $E_t[v(g(t), t)]$ subject to IR_Y and $EPIC$, while $g^Y(c)$ maximizes $E_t[v(g(N), t)]$ subject to IR_N and MH_N since, without loss of optimality $g^Y(c)(N) = 0$ and $g^N(c)(t) = 0$ for all $t \in T$.

I will show that $V(g^Y(c)) \equiv E_t[v(g^Y(c)(t),t)]$ is continuous and decreasing in c, while $V(g^N(c)) \equiv E_t[v(g^N(c)(t),t)]$ is continuous and increasing in c.

For the monotonicity results, note that the set of contracts that satisfy IR_N is independent of c and the set of contracts that satisfy MH_N is expanding in c. It follows that $V(g^N(c))$ is (weakly) increasing in c. The set of vectors of contracts that satisfy IR_Y is shrinking in c while the set of vectors of contracts that satisfy $EPIC$ is independent of c. It follows that $V(g^Y(c))$ is weakly decreasing in c.

For continuity, note that in both cases, objective function is continuous in the vector $((g(t))_{t\in T}, g(N))$. The vector of contracts $g^Y(c)$ maximizes $E_t[v(g^Y(c)(t),t)]$ subject to $g^Y(c) \in \Gamma^Y(c)$, where $\Gamma^Y(c)$ is the set of vectors of contracts that satisfy IR_Y and $EPIC$. This set is compact valued, as we have seen, as well as being upper- and lower-hemicontinuous. To see that $\Gamma^Y(c)$ is upper-hemicontinuous let $c^k \to c$ and $g^k \in \Gamma^Y(c^k)$ with $g^k \to g$. We can take each $g^Y(c)(t)$ to be in a bounded set, so that $E_t[u(g(t),t)]$ is continuous in $g(t)$, by the dominated convergence theorem. Therefore, for any ϵ and large k:

$$E_t[u(g(t),t)] - c \geq E_t[u(g^k(t),t)] - \epsilon - c \geq$$
$$E_t[u(g^k(t),t)] - c^k - 2\epsilon \geq E_t[u(0,t)] - 2\epsilon.$$

Taking ϵ to zero, we see that $E_t[u(g(t),t)] - c \geq E_t[u(0,t)]$. Similarly, $g(t)$ satisfies $EPIR$ for all t, so that Γ^Y is upper hemicontinuous.

To see that Γ^Y is lower hemicontinuous, $c^k \to c$ and $g \in \Gamma^Y(c)$. Then let $g^m(t) = (1-\delta_m)g(t) + \delta_m(\hat{g}(t))$, where $u(\hat{g}(t),t) > u(g(t),t)$ and $\delta_m \to 0$. Since $u(a,p,t)$ is strictly decreasing in p for all $t \in T$, such a $\hat{g}(t)$ always exists. Note that, given that IR_Y and $EPIR$ are satisfied at c by $g(t)$, they are strictly satisfied by $g^m(t) = (1 - \delta_m)g(t) + \delta_m(\hat{g}(t))$ for c^m sufficiently close to c. Moreover, $g^m(t) \to g$. Selecting a subsequence, c^m of the sequence c^k for which g^m satisfies the constraints at c^m, we see that $\Gamma^Y(c)$ is lower semicontinuous.

A similar argument, mutatis mutandis, shows that $\Gamma^N(c)$ – the set of contracts that satisfy IR_N and MH – is a continuous correspondence in c.

We have seen in the proof of Lemma 1 that both $\Gamma^Y(c)$ and $\Gamma^N(c)$ can be taken to be compact valued for all c. It follows by Berge's maximum theorem that $V^Y(c)$ and $V^N(c)$ are continuous in c.

Now note that at $c = 0$ any contract that is feasible without evidence is also feasible with evidence, so that $V(g^Y(0)) \geq V(g^N(0))$. For sufficiently high c, by assumption 5, any feasible contract with evidence does worse than the outside option. It follows that $V(g^Y(c)) < V(g^N(c))$ for c sufficiently large. By the intermediate value theorem it follows that there exists a $c^* \geq 0$ such that $V(g^Y(c)) \geq V(g^N(c))$ if and only if $c \leq c^*$.

Now let $c^{**} = max\{c^*, \tilde{c}\}$. If $c > c^{**}$ then the optimal contract without evidence is optimal overall, and this contract must be $(\alpha^*, \hat{p}(\alpha^*))$. At $c > c^{**} \geq \tilde{c}$ MH_N is not binding for this contract. This completes the proof of Lemma 2.

Proof of Proposition 3: Suppose that $E_t[v(g(t)(\tilde{c}), t)] \geq E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$. Note that $E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$ is the best possible expected payoff to the principal from a contract that does not induce evidence acquisition subject to IR_N and ignoring MH. If $c < \tilde{c}$ then $(\alpha^*, \hat{p}(\alpha^*)$ is not feasible and $E_t[v(g(t)(\tilde{c}), t)] \geq E_t[v(\alpha^*, \hat{p}(\alpha^*), t)] > E_t[v(g(N), t)]$ for any feasible contract $g(N)$ that does not induce evidence acquisition. Therefore, in this case, the principal chooses not to induce evidence acquistion only for $c \geq \tilde{c}$. It follows from Lemma 2 that for c less than some $c^{**} \geq \tilde{c}$ the optimal contract induces evidence acquisition and has IR_Y binding, while for $c > c^{**}$ the contract is $(\alpha^*, \hat{p}(\alpha^*))$. Since IR_Y becomes stricter as c increases, the principal's payoff is decreasing for $c < c^{**}$ and constant for $c \geq c^{**}$. Therefore the principal's payoff in non-increasing, as claimed.

Now suppose that $E_t[v(g(t)(\tilde{c}), t)] < E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$. Then at $\tilde{c}$ the principal strictly prefers to induce evidence acquisition. By the continuity of the principal's value functions, established in the proof of Lemma 2, there is some interval $[c^*, \tilde{c})$ in which the principal chooses to not to induce evidence acquisition and MH is binding. Since MH becomes easier to satisfy as c increases, it follows that the principal's payoff is increasing on this interval. By Lemma 2, $c^* > 0$ and the IR_Y is binding for $c < c^*$so that the principal's profit is decreasing on $[0, c^*)$. It follows that the principal's profit is U-shaped in cost, as claimed.

Finally, suppose that the value of information is non-positive and $u(\alpha^*, \hat{p}(\alpha^*), t) < E_t[u(\alpha^*, \hat{p}(\alpha^*))]$ for some $t \in T$. Then:

$$E_t[v(g(t)(\tilde{c}), t)] < max_{(g(t))_{t\in T}: u(g(t),t) \geq u(0,t)\ \forall t} E_t[v(g(t), t)] \leq E_t[v(\alpha^*, \hat{p}(\alpha^*), t)]$$

As long as $\tilde{c} > 0$. The hypothesis that $u(\alpha^*, \hat{p}(\alpha^*), t) - u(0, t) > 0$ for some $\tilde{t} \in T$ insures that this is the case, since then

$$E_t[u(\alpha^*, \hat{p}(\alpha^*)] - E_t[u(0, t)] = 0 < E_t[\max\{u(0, t), u(\alpha^*, \hat{p}(\alpha^*), t)\}\mathbf{1}\{t = \tilde{t}\}] - c$$

$$\leq E_t[max\{u(0, t), u(\alpha^*, \hat{p}(\alpha^*), t)\}]$$

as long as c is sufficiently small, so that MH is sometimes violated by $(\alpha^*, \hat{p}(\alpha^*))$.

5.2 Omitted proofs from Section 3

Proof of Lemma 9: In the main text we have established that the only candidate solutions have $a(L) = a(H) = 1$, with the pair of premiums either $(p_1(L), p_2(H)) \equiv (W - U^{-1}(\tilde{u}^L + \frac{c}{1-\lambda})), W - U^{-1}(\tilde{u}^H))$ or $(p_2(L), p_2(H)) \equiv (W - U^{-1}(\tilde{u}^L), W - U^{-1}(\tilde{u}^H + \frac{c}{\lambda})$. We will show that the latter implies a higher payoff for the insurer.

In order to obtain a contradiction, suppose not. Then it must be that $\lambda p_1(H) + (1-\lambda)p_1(L) > \lambda p_2(H) + (1-\lambda)p_2(H)$. This implies that $\lambda(W - p_1(H)) + (1-\lambda)(W - p_1(H)) < \lambda(W - p_2(H)) + (1-\lambda)(W - p_2(H))$. Note that $W - p_1(H) < W - p_2(H)$ and $W - p_1(L) > W - p_2(L)$. This implies that we can find a positive constant, δ such that $(W - p_1(L) + \delta, W - p_1(H))$, viewed as a lottery over consumption levels, is a mean preserving spread of $(W - p_2(L), W - p_2(H))$.
Since the consumer is risk averse, he must therefore prefer $(W - p_2(L), W - p_2(H))$ to $(W - p_1(L) + \delta, W - p_1(H))$, which is clearly prefered to $(W - p_1(L), W - p_1(H))$. However, by construction, both contracts give the same ex-ante expected utility, so that the consumer is indifferent between them. This is a contradiction. It follows that the initial hypothesis is false, and the pair of premiums $(p_2(L), p_2(H))$ gives a higher profit to the insurer.

5.3 Analysis of the unit valuation monopoly model

Here we consider the special case of the general model with $A = [0, 1]$ and

$$u(a, p, t) = t \cdot a - p$$

and

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

Notice that $a^* = a^*_t = 1$ for all $t \in T$. It follows that in this case the value of information is zero. We know from the general results that when evidence is acquired on-path the contract offered when t is disclosed is $(1, t - c)$, giving zero consumer surplus and profit equal to:

$$V^E(c) = E[t] - c.$$

When evidence is not acquired on-path, the best contract subject to IR_N is $(1, E[t])$. This is feasible, and hence optimal, if c is higher than $\tilde{c}$ satisfying:

$$\tilde{c} = \pi[t \geq E[t]](E[t|t \geq t] - E[t]).$$

Now suppose that $c < \tilde{c}$. Now, for $c < \tilde{c}$, let (a, p) be any feasible contract satisfying MH with $a < 1$. That is:

$$aE[t] - p \geq (aE[t|a \cdot t \geq p] - p)\pi[a \cdot t \geq p] - c.$$

Consider $a' > a$ and $p' = p$. We see that the left hand side is increased by $(a' - a)E[t]$, while the right hand side is increased by at most $(a' - a)E[t|a \cdot t \geq p]\pi[a \cdot t \geq p]$ (since $\{t : a' \cdot t \geq p\}$ is no larger). The difference between the increase in the left hand side and the increase in the right hand side is therefore at most:

$$(a' - a)E[t|a \cdot t < p]\pi[a \cdot t < p] > 0.$$

Under the assumption that $a \cdot t < p$ for some t. If not, let $\underline{t}$ be the lowest type and set the new contract at $(1, \underline{t})$, which must be feasible and increase profit.

It follows that MH is slack at (a', p) (as must be IR_N) hence we can increase the price by some $\epsilon > 0$, increasing profit.

Therefore, $a < 1$ is never optimal. Since we know from the general results that MH always binds in the optimal mechanism without evidence, when $c < \tilde{c}$. The optimal mechanism must be $(1, p(c))$, where $p(c)$ satisfies:

$$E[t] - p(c) = (E[t|t \geq p(c)] - p(c))\pi[t \geq p(c)] - c.$$

It follows that $V(c)$ satisfies:

$$V(c) = E[t|t < V(c)] + \frac{c}{\pi[t < V(c)]}$$

Which must be increasing in c. Conversely, since aggregate welfare is constant in the interval $(c^*, \tilde{c})$ (since the allocation is constant and no evidence is acquired), the consumer surplus in decreasing in c. Since the value of information in this environment is zero, while $E_t[max\{u(0,t), u(a^*,t) - E_t[u(a^*,t)]\}] > 0$, Proposition 1 implies that $c^* < \tilde{c}$.

These features are depicted in Figure 2.