

Robust contracting under common value uncertainty

SARAH AUSTER

Department of Decision Sciences and IGIER, Bocconi University

A buyer makes an offer to a privately informed seller for a good of uncertain quality. Quality determines both the seller's valuation and the buyer's valuation, and the buyer evaluates each contract according to its worst-case performance over a set of probability distributions. This paper demonstrates that the contract that maximizes the minimum payoff over *all possible probability distributions* of quality is a screening menu that separates all types, whereas the optimal contract for *any given probability distribution* is a posted price, which induces bunching. Using the ε -contamination model, according to which the buyer's utility is a weighted average of his single prior expected utility and the worst-case scenario, the analysis further shows that for intermediate degrees of confidence, the optimal mechanism combines features of both of these contracts.

KEYWORDS. Ambiguity, optimal contracting, lemons problem.

JEL CLASSIFICATION. D81, D82, D86.

1. INTRODUCTION

The optimal design of contracts in the presence of asymmetric information has been the subject of investigation for several decades and its theoretical analysis has generated an array of powerful results. Most of this literature adopts the subjective expected utility model according to which contracting parties have a single subjective prior belief about the fundamentals. In real-life contracting situations, e.g., when buying a house or when investing in a foreign country, the involved parties rarely have a precise idea about the underlying probability distribution, either because they do not have enough experience or because they do not have sufficient information. It is well established that a lack of knowledge about the probability distribution can have important behavioral implications that are incompatible with the subjective expected utility hypothesis. This gives rise to the question of how the presence of uncertainty over probabilistic scenarios affects the optimal design of contracts and the implemented allocation. This paper analyzes a bilateral trade model that allows for common values—an important

Sarah Auster: sarah.auster@unibocconi.it

This work was undertaken as part of my PhD thesis at the European University Institute. It previously circulated under the title “Bilateral Trade Under Ambiguity.” I am very grateful to my supervisor Piero Gottardi. I would also like to thank Arpad Abraham, Subir Bose, Itzhak Gilboa, Emeric Henry, Peter Klibanoff, Frederic Koessler, Nenad Kos, David Levine, Massimo Marinacci, Matthias Messner, Sujoy Mukerji, Andrew Postlewaite, David Pothier, Tomasz Strzalecki, Balazs Szentes, and the participants of various seminars for very helpful comments.

feature in many real-life contracting situations—and an ambiguous trading environment.

In the environment considered, there is a risk-neutral buyer (she) who makes an offer to a risk-neutral seller (he), who is privately informed about the quality of his good. The paper first presents an introductory example in which quality is either high or low and then studies the richer environment in which quality belongs to an interval of values. In contrast to the classic setting, it is assumed that the buyer has ambiguous beliefs about the distribution of quality, which determines the valuation of both trading parties. The buyer's preferences are represented by the maxmin expected utility model (Gilboa and Schmeidler 1989), according to which the buyer evaluates her choices with the most pessimistic probability distribution in a set of distributions.

For the case where this set is a singleton, Samuelson (1984) shows that the optimal mechanism is a take-it-or-leave-it price. The results of this paper demonstrate that if the extent of ambiguity the buyer faces is sufficiently large and if values are interdependent, the optimal mechanism is a screening menu rather than a posted price. To show this, the paper first studies the case of Knightian uncertainty, where the buyer considers the worst-case payoff over all possible probability distributions of quality. Here the buyer optimally proposes a contract that equalizes her payoff across all seller types and thereby hedges against the ambiguity she perceives. This contract can be interpreted as maximally robust, since it yields the same expected payoff across all possible probabilistic scenarios. The analysis shows that the nature of the robust mechanism crucially depends on the relation between the buyer's and the seller's valuations: if the buyer's valuation strictly increases in the seller's type, the optimal mechanism is a screening menu that perfectly separates all seller types, whereas if the buyer's and seller's valuations are independent, the maximally robust mechanism is the pooling price. The intuition is that a separating menu allows the price of the good to increase with the buyer's valuation of the good, thereby balancing her payoff across the different realizations of her and the seller's valuation.

The paper also studies the case in which the buyer's ignorance is less extreme. To parameterize the buyer's demand for robustness, her preferences are represented by the ε -contamination model, a special case of the maxmin expected utility model, which nests the case of Knightian uncertainty on the one hand and the case of a single subjective prior belief on the other hand. According to this representation, the buyer has a reference distribution but entertains some doubt regarding that distribution, captured by the parameter ε . The buyer's confidence in the model distribution determines the nature of the optimal mechanism. If ε is sufficiently small, the optimal mechanism is a posted price as in Samuelson (1984), whereas if ε is sufficiently large, the optimal mechanism is a screening menu that perfectly hedges against ambiguity. For intermediate values of ε , the optimal contract combines features of both of these mechanisms: low quality sellers are bunched at a base price, while high quality sellers are separated by the mechanism. This hybrid contract solves the trade-off between maximizing the buyer's expected utility evaluated at the reference distribution and limiting her minimal payoff in the worst-case scenario. In particular, by screening sellers with a valuation above the

base price, the buyer avoids the possibility of trading with probability 0 in some states of the world.

Given the longstanding debate on the economic implications of ambiguity as opposed to risk, I discuss how the optimal contract in the proposed setting differs from the benchmark model in which the buyer is ambiguity neutral but risk averse. Although the optimal contract under risk aversion may also be a separating menu, the cases in which it coincides with the optimal mechanism under ambiguity aversion are nongeneric. Moreover, the discussion demonstrates that there are situations in which a higher degree of risk aversion makes the pooling price optimal, while a higher degree of ambiguity aversion generally favors separation.

Finally, in some applications of the maxmin expected utility model, results crucially rely on the feature that preferences are kinked (e.g., [Dow and Ribeiro da Costa Werlang 1992](#); [Condie and Ganguli 2017](#)). This is not the case in the model studied here because the optimality of separating menus in the considered environment is driven by the buyer's desire to hedge against ambiguity, a feature of all decision models that capture ambiguity averse behavior. This is illustrated by extending the characterization of the optimal contract for the binary type case to the smooth ambiguity model, introduced by [Klibanoff et al. \(2005\)](#). The characterization shows that the solution of the buyer's optimization problem under smooth ambiguity aversion is a convex combination of the solution under maxmin expected utility and ambiguity neutrality.

Related literature

This paper is part of a growing literature on robust contracting in an uncertain environment, which includes work on procurement contracts ([Garrett 2014](#)), optimal delegation mechanisms ([Frankel 2014](#); [Carrasco and Moreira 2013](#)), and optimal incentive contracts in the presence of moral hazard ([Carroll 2015](#); [Carroll and Meng 2016](#); [Antić 2014](#)). Crucially, and in contrast to most prior work in mechanism design, the principal in these models evaluates contracts according to their worst-case performance, e.g., over the agent's preferences or over the set of available technologies.

There is also a small number of papers that study the maxmin optimal contract in the canonical principal–agent problem with hidden information, as this one does. In contrast to this work, existing papers assume that the principal's and agent's valuations are independent. [Bergemann and Schlag \(2011\)](#) show that under the assumption of independent private values, the maxmin optimal mechanism is a posted price. Intuitively, if the principal only faces ambiguity over the agent's acceptance decision, the worst-case probabilistic scenario is always the one that maximizes the probability that the agent rejects. This implies that the optimization problem under maxmin preferences is equivalent to the optimization problem under a single pessimistic prior and [Samuelson's \(1984\)](#) result applies. [Bergemann and Schlag \(2011\)](#) also consider the mechanism that minimizes maximal regret. In contrast to maxmin expected utility, the minimax regret criterion generates a regret trade-off that makes randomization across prices optimal. In recent work, also [Carrasco et al. \(2017\)](#) consider a monopoly pricing

model with independent private values and a principal with maxmin preferences. In their environment, the principal has partial probabilistic information about the agent's valuation, such as the mean or the variance, which makes a random pricing rule optimal.

There exists a complementary line of literature that studies mechanism design problems in which the agent rather than the principal faces uncertainty about the underlying probabilistic environment. This literature includes the work of [Bose and Mutuswami \(2012\)](#) and [Wolitzky \(2016\)](#), both of which investigate the implementability of efficient trade in the canonical [Myerson and Satterthwaite \(1983\)](#) environment with the assumption that agents perceive uncertainty about the probability distribution of the opponent's type. Also [Bose et al. \(2006\)](#), [Bose and Daripa \(2009\)](#), and [Bodoh-Creed \(2012\)](#) introduce ambiguity on the agent's side, but in contrast to the work mentioned before, the focus of these papers lies on revenue maximization. Finally, [Bose and Renou \(2014\)](#) and [di Tillio et al. \(2017\)](#) show that the designer may benefit from introducing ambiguity via the mechanism.

The rest of the paper is organized as follows. [Section 2](#) presents the introductory example, which demonstrates the main features of robust contracting in the considered environment. [Section 3](#) then introduces the main model. I first characterize the optimal contract for the case in which the buyer considers the worst-case scenario over all possible probability distributions. Next, I study the case of moderate ambiguity and show how the features of the optimal contract depend on the buyer's attitude toward the ambiguity she faces. [Section 4](#) discusses the differences to risk aversion and explains how the findings of the model extend when preferences are smooth rather than kinked. [Section 5](#) concludes.

2. INTRODUCTORY EXAMPLE

A risk-neutral buyer makes an offer to a risk-neutral seller who possesses one unit of an indivisible good. The seller is privately informed about the quality of the object, which can be either high or low. Quality determines both the seller's and the buyer's valuations, implying that the seller knows both his own valuation $c \in \{c_l, c_h\}$ as well as the buyer's valuation $v \in \{v_l, v_h\}$, whereas the buyer knows neither of these values. I assume that both the seller and the buyer value high quality more than low quality, i.e., $c_h > c_l$, $v_h > v_l$, and that the buyer's value always exceeds the seller's value: $v_i > c_i$, $i = l, h$. The buyer proposes a menu of contracts, $\{(x(c), t(c))\}_{c \in \{c_l, c_h\}}$, consisting of a trading probability $x(c)$ and a transfer $t(c)$ for each type of good. If the seller reveals truthfully his type, the buyer's and the seller's payoffs as a function of the seller's type are given by $\pi_b(c) = x(c)v(c) - t(c)$ and $\pi_s(c) = t(c) - x(c)c$, respectively.

The buyer faces ambiguity over the quality distribution and has maxmin expected utility preferences ([Gilboa and Schmeidler 1989](#)). Under this representation, a decision maker evaluates her choices with the worst probability distribution in a convex set of distributions. Letting σ denote the probability that the quality of the object is high, this amounts to the buyer minimizing over an interval of values of σ . Let E_σ

denote the expectation operator with respect σ . The buyer's payoff function is then given by

$$\inf_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} E_{\sigma}[\pi_b(c)].$$

The buyer maximizes this payoff function subject to the incentive and individual rationality constraints of each type of seller:

$$\begin{aligned} t(c_l) - x(c_l)c_l &\geq t(c_h) - x(c_h)c_l, \\ t(c_h) - x(c_h)c_h &\geq t(c_l) - x(c_l)c_h, \\ t(c_i) - x(c_i)c_i &\geq 0, \quad i = l, h. \end{aligned}$$

Since the seller knows his type, the constraints of the buyer's optimization problem are not affected by the presence of ambiguity in this environment. This, and the fact that the buyer's objective is weakly decreasing in $t(c_l)$ and $t(c_h)$, implies that the solution to the buyer's optimization problem satisfies some well established properties (see, for example, [Salanié 2005](#), Chapter 2): the incentive compatibility constraint of the low type seller and the individual rationality constraint of the high type seller are binding, while the remaining constraints are slack. Furthermore, the low type seller trades with probability 1. With these properties, the menu of contracts is completely characterized by the trading probability of the high type seller $x(c_h)$. For notational convenience, let this probability be denoted by α :

$$\{(x(c_l), t(c_l)), (x(c_h), t(c_h))\} = \{(1, \alpha c_h + (1 - \alpha)c_l), (\alpha, \alpha c_h)\}.$$

REMARK. Note that there are two alternative interpretations of the model. In the interpretation followed throughout the paper, the good is indivisible and α is the probability of trade. In an alternative interpretation, the seller possesses one unit of a perfectly divisible good, utility functions are multiplicatively linear, and α is a quantity. Under this interpretation, the menu is a nonlinear pricing schedule with a quantity discount. Low quality is traded in large quantity at a low price, while high quality is traded in small quantity at a high price.

Given the properties stated above, the buyer's payoff can be stated as a function of α and her optimization problem becomes

$$\max_{\alpha \in [0, 1]} \inf_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} \{\sigma \alpha (v_h - c_h) + (1 - \sigma)(v_l - \alpha c_h + (1 - \alpha)c_l)\}.$$

To derive the solution to this problem, consider first the case in which the set of probability distributions is a singleton, so that the buyer is a subjective expected utility maximizer with a single prior σ . [Samuelson \(1984\)](#) shows that the buyer's optimal contract is a posted price, which the seller either accepts or rejects. In the two-type setting, this result can easily be seen by considering the buyer's objective function. Since $E_{\sigma}[\pi_b]$ is linear in α , the optimization problem has a corner solution. The mechanism characterized by $\alpha = 0$ is a separating price equal to c_l , while that characterized by $\alpha = 1$ is the

pooling price equal to c_h . Pooling is optimal if the probability that the seller's type is high is large enough, which is the case if

$$\sigma v_h + (1 - \sigma)v_l - c_h \geq (1 - \sigma)(v_l - c_l) \quad \text{or, equivalently,} \quad \sigma \geq \frac{c_h - c_l}{v_h - c_l}.$$

Proposition 2.1 shows that if the set of probability distributions is not a singleton, the buyer's optimal contract is a posted price if and only if this price is optimal under all probability distributions in the interval $[\underline{\sigma}, \bar{\sigma}]$. Otherwise, the buyer optimally offers a separating menu.

PROPOSITION 2.1. *Let $\tilde{\sigma} := \frac{c_h - c_l}{v_h - c_l}$. The buyer optimally proposes a menu characterized by*

$$\alpha^* = \begin{cases} 0 & \text{if } \bar{\sigma} \leq \tilde{\sigma}, \\ 1 & \text{if } \underline{\sigma} \geq \tilde{\sigma}, \\ \frac{v_l - c_l}{v_h - c_l} & \text{otherwise.} \end{cases}$$

See [Appendix A](#) for the proofs of most of the propositions.

The parameter $\tilde{\sigma}$ is the subjective prior under which all values of α yield the same payoff. If $\underline{\sigma} \geq \tilde{\sigma}$, the optimal value of α is equal to 1, since the contract that maximizes the buyer's expected payoff is the pooling price c_h for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. Similarly, if $\bar{\sigma} \leq \tilde{\sigma}$, the optimal value of α is equal to 0, because the separating price c_l is optimal for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$. Ambiguity aversion thus affects the buyer's utility but not her choice of contract. If $\underline{\sigma} < \tilde{\sigma} < \bar{\sigma}$, the buyer optimally proposes a separating menu, characterized by $\alpha^* = \frac{v_l - c_l}{v_h - c_l}$. This value of α is the contracting parameter under which the buyer's payoff conditional on the seller's type being high equals her payoff conditional on the seller's type being low. Intuitively, offering a separating menu rather than a posted price allows the buyer to limit the extent of ambiguity over the probability with which the seller accepts and the net gain when trade occurs, thereby balancing her payoff across the different states of the world. The menu characterized by $\alpha^* = \frac{v_l - c_l}{v_h - c_l}$ can be viewed as a robust contract because it makes the buyer's payoff independent of the underlying type distribution and thereby yields a "safe" payoff equal to $\frac{(v_h - c_h)(v_l - c_l)}{v_h - c_l}$. Offering a separating menu thus hedges against ambiguity in this environment. Hedging is optimal if and only if the environment is sufficiently ambiguous and the buyer is sufficiently ambiguity averse, i.e., if $[\underline{\sigma}, \bar{\sigma}]$ is large enough.

This is illustrated in [Figure 1](#). Under the assumption $\underline{\sigma} < \tilde{\sigma} < \bar{\sigma}$, the expected payoff of a buyer with subjective prior $\underline{\sigma}$ is downward sloping in α (solid line), while the expected payoff of a buyer with subjective prior $\bar{\sigma}$ is upward sloping in α (dashed line). All expected payoff functions $E_\sigma[\pi_b]$, $\sigma \in (\underline{\sigma}, \bar{\sigma})$ lie in between these two benchmarks and intersect at $\alpha^* = \frac{v_l - c_l}{v_h - c_l}$. For $\alpha < \frac{v_l - c_l}{v_h - c_l}$, the worst case for the buyer is that the probability with which the seller has a high quality good is large, whereas for $\alpha > \frac{v_l - c_l}{v_h - c_l}$, the worst case is that the probability of this event is small. Since the buyer's expected payoff evaluated at $\bar{\sigma}$ is upward sloping in α , while her expected payoff evaluated at $\underline{\sigma}$ is downward sloping in α , her maxmin expected payoff (thick curve) is maximized at $\alpha^* = \frac{v_l - c_l}{v_h - c_l}$.




FIGURE 1. Expected payoffs $E_\sigma[\pi_b]$ for $\sigma \in [\underline{\sigma}, \overline{\sigma}]$.

3. THE MAIN MODEL

Consider now the case where the seller’s type takes a value in the interval $[0, 1]$. As before, the seller knows his own valuation for the object c and also the buyer’s valuation v , whereas the buyer is uncertain about these values. Let the differentiable function $v(c)$ define the relation between the two values, and assume $v'(c) \geq 0$ and $v(c) > c$ for all $c \in [0, 1]$. That is, both the seller’s and the buyer’s value for the object are increasing in its quality, and the buyer’s value strictly exceeds the seller’s value.

In this environment a menu of contracts is defined as $\{(x(c), t(c))\}_{c \in [0, 1]}$, and the incentive and individual rationality constraints of the seller are given by

$$t(c) - x(c)c \geq t(\tilde{c}) - x(\tilde{c})c \quad \forall c, \tilde{c},$$

$$t(c) - x(c)c \geq 0 \quad \forall c.$$

Samuelson (1984, p. 997) shows that this set of constraints implies that $x(c)$ weakly decreases in c and that $\pi'_s(c) = -x(c)$ almost everywhere. With these properties, the buyer’s and seller’s payoffs can be derived as a function of $x(c)$ only. Condition $\pi'_s(c) = -x(c)$ together with $\pi_s(1) = 0$ yields

$$\pi_s(c) = \int_c^1 x(u) \, du.$$

The term $\int_c^1 x(u) \, du$ is the information rent paid to type c . The buyer’s payoff as a function of the seller’s type, $\pi_b(c)$, is then given by the difference between the expected value of the realized gains from trade and the information rent paid to the seller. That is,

$$\pi_b(c) = x(c)\Delta(c) - \int_c^1 x(u) \, du,$$

where $\Delta(c) \equiv v(c) - c$ denotes the gains from trade when the seller’s type is c .

3.1 Unique subjective prior

Consider first the benchmark case, where the buyer has a single prior belief F . This specification corresponds to the standard Bayesian setting, for which Samuelson (1984) shows that the optimal mechanism for the buyer is a posted price. Assume that F is twice differentiable and let $f(c) = F'(c)$ denote the density function. Defining Ψ as the set of nonincreasing functions $x : [0, 1] \rightarrow [0, 1]$, the buyer solves the problem

$$\max_{x \in \Psi} E_F[\pi_b(c)] = \int_0^1 \left(x(c)\Delta(c) - \int_c^1 x(u) du \right) f(c) dc. \quad (I)$$

After an integration by parts, the buyer's expected payoff can be written as $E_F[\pi_b(c)] = \int_0^1 (\Delta(c)f(c) - F(c))x(c) dc$, where the function $H(c) \equiv \Delta(c)f(c) - F(c)$ captures the marginal benefit of increasing $x(c)$. If $H(c)$ is strictly decreasing in c , the optimal mechanism is characterized by the step function

$$x^*(c) = \begin{cases} 1 & \text{if } c \leq c^*, \\ 0 & \text{if } c > c^*, \end{cases}$$

where c^* is such that $H(c^*) = 0$ if $H(1) < 0$ and $c^* = 1$ otherwise. This mechanism corresponds to a posted price equal to c^* . The buyer's payoff associated to the posted price varies with the seller's type c . In particular, the minimum payoff is given by $\bar{\pi}^{\min} \equiv \{\Delta(0) - c^*, 0\}$, which is obtained either when the seller is of the lowest type so that the buyer's valuation is $\Delta(0)$ or when the seller rejects the buyer's offer.

3.2 Knightian uncertainty

Consider now the polar case, where the buyer evaluates possible contracts by their worst-case performance over all probability distributions on $[0, 1]$. This case of Knightian uncertainty has received a lot of attention in the literature on robust contracting, e.g., Frankel (2014), Garrett (2014), and Carroll (2015). Under this specification, the buyer solves the problem

$$\max_{x \in \Psi} \inf_{c \in [0, 1]} \left\{ x(c)\Delta(c) - \int_c^1 x(u) du \right\}. \quad (II)$$

The following proposition characterizes the solution to this problem.

PROPOSITION 3.1. *The unique solution of problem (II) is given by*

$$x^*(c) = \exp \left[- \int_0^c \frac{v'(t)}{\Delta(t)} dt \right] \quad \text{for all } c \in [0, 1].$$

The optimal mechanism for a buyer who minimizes over all types $c \in [0, 1]$ is a mechanism that equalizes her payoff across all types. Noting that $\pi'_b(c) = x'(c)\Delta(c) + x(c)v'(c)$, the mechanism that yields a constant payoff across c solves the differential equation

$$x'(c)\Delta(c) + x(c)v'(c) = 0$$

with $x(0) = 1$. The solution to this differential equation is given by the function $x^*(c)$, as defined in Proposition 3.1. Since x^* is nonincreasing, the monotonicity constraint for incentive compatibility is automatically satisfied. The mechanism characterized by x^* can be viewed as the analogue to the contract characterized by $\alpha^* = \frac{v_l - c_l}{v_h - c_l}$ in the binary case. Its key feature is that it yields a payoff that does not depend on the underlying probability distribution over types. Instead, the buyer obtains a safe payoff equal to

$$\bar{\pi}^{\max} \equiv \Delta(1) \exp \left[- \int_0^1 \frac{v'(t)}{\Delta(t)} dt \right].$$

It should be emphasized that the nature of the mechanism characterized by $x^*(c)$ crucially depends on the relation between the buyer's valuation and the seller's valuation. Noting that the first derivative of the optimal trading probability is given by

$$x^{*'}(c) = - \exp \left[- \int_0^c \frac{v'(t)}{\Delta(t)} dt \right] \frac{v'(c)}{\Delta(c)},$$

it is easy to see that $x^*(c)$ strictly decreases in c if and only if the buyer's valuation $v(c)$ strictly increases in c . In particular, if there are subsets of $[0, 1]$ for which $v'(c) = 0$, the buyer optimally bunches types on those subsets but separates the rest. The intuition is that a strictly increasing $v(c)$ favors trading with high types and thus calls for an asymmetric treatment of sellers so as to equalize the buyer's payoff across types. Vice versa, if $v(c)$ is constant on some part of the domain, equalizing the buyer's payoff across sellers requires bunching seller types in that part of the domain.

In the independent private value case, where $v(c) = \bar{v}$ for all $c \in [0, 1]$, the optimal mechanism is, in fact, characterized by $x^*(c) = 1$ for all c . This mechanism corresponds to the pooling price, which yields a safe payoff equal to $\bar{v} - 1$. Once we move away from the independent private value case, the optimal mechanism is no longer the pooling price—or any other posted price—but a separating menu. To gain some intuition, note that if the buyer offers the pooling price and $v'(c) > 0$ on some part of the domain, the buyer obtains her minimum payoff when trading with the lowest type. This minimum payoff can be increased by reducing the trading probability of higher types so as to reduce the information rent paid to the lowest type. Under the optimal contract instead, the trading probability strictly decreases on the parts of the domain where v strictly increases. This allows the buyer to offset the additional information rent paid to lower types with a higher probability of trade, thereby balancing the buyer's payoff across seller types. In particular, if v is strictly increasing on the whole domain, the maximally robust mechanism is a screening menu that separates all seller types.

REMARK. Note that if instead $v'(c) \leq 0$ for all $c \in [0, 1]$, the problem is less interesting. Here $\pi_b(c)$ necessarily decreases in c , implying that the worst-case scenario always includes the degenerate distribution with a mass point at $c = 1$. The buyer thus optimally proposes the pooling price. Also when the assumption of strictly positive gains from trade is violated, the contracting problem becomes trivial. In particular, if there exists some $\tilde{c} \in [0, 1]$ such that $\Delta(\tilde{c}) \leq 0$, not trading is a weakly dominant strategy for the buyer: if the seller's type is \tilde{c} with probability 1, there is no contract that can yield a strictly positive payoff for the buyer.

It is interesting to point out that whenever the gains from trade are constant across quality so that $\Delta(c) = \delta$, the mechanism that maximizes (II) is characterized by

$$x^*(c) = e^{-\frac{c}{\delta}} \quad \text{for all } c \in [0, 1].$$

The trading probability is thus exponentially distributed, scaled by a factor δ . One feature of the exponential probability distribution is that the density function and the decumulative distribution function decrease proportionally so that the hazard rate is constant.¹ For the optimal trading probability function $x^*(c)$, this implies that the expected value of the realized gains from trade, $x^*(c)\delta$, and the information rent, $\int_c^1 x^*(u) du$, change at the same rate, thereby yielding a constant payoff in c . Clearly, the optimality of the exponential probability distribution in the problem under consideration hinges on the assumption that gains from trade are indeed constant. However, the result points to a more general property of the optimal mechanism, namely an inverse relation between the gains from trade $\Delta(c)$ and the hazard rate of $x^*(c)$: if gains from trade are increasing in c , $x^*(c)$ has a decreasing hazard rate, whereas if gains from trade are decreasing in c , $x^*(c)$ has an increasing hazard rate.

3.3 Intermediate case

The maximally robust mechanism characterized in Proposition 3.1 provides a useful benchmark. Nevertheless, the Knightian decision criterion can be viewed as somewhat extreme. Given the stark difference between the optimal mechanism under a single subjective prior and the mechanism characterized in Proposition 3.1, it is then of interest to analyze the case where the buyer's demand for robustness is less extreme. To make this analysis tractable, the buyer's preferences will be represented by the ε -contamination model, axiomatized by Kopylov (2016) among others, according to which the buyer has some reference probability distribution F in mind but considers an ε perturbation around it. Under this representation, the buyer's utility function is given by $(1 - \varepsilon)E_F[\pi_b(c)] + \varepsilon \inf_{G \in \text{cl}(\text{co}\Delta)} E_G[\pi_b(c)]$, $\varepsilon \in [0, 1]$, where F belongs to the closed convex hull of Δ , which is an exogenous information set known to the decision maker that contains the true probability law on the state space.

In what follows I assume that the buyer has no exogenous information about the true distribution of seller types so that Δ contains all probability measures on $[0, 1]$. This assumption has the implication that the buyer's perceived ambiguity is uniform across the states of the world, which is not without loss of generality. For example, for the case of independent private values, Carrasco et al. (2017) show that if ε is equal to 1 and Δ includes only those probability distributions with the same mean, the optimal mechanism is a random pricing rule, whereas Proposition 3.1 shows that if Δ contains all probability distributions on $[0, 1]$, the optimal mechanism is the (deterministic) pooling price. How the optimal mechanism changes with Δ when it takes more general forms might pose an interesting question for future research.

¹Letting g and G , respectively, denote the density function and the cumulative distribution function, the hazard rate is defined by $\frac{g(t)}{1-G(t)}$.

Given this restriction, the buyer’s payoff function is given by the weighted average of her subjective expected utility evaluated at F and the minimum payoff over all state realizations. The buyer’s decision model is thus as if she thinks that with probability $1 - \varepsilon$, the seller’s type is distributed according to F , while with the complementary probability it could be any other probability distribution on $[0, 1]$. The confidence parameter ε measures the buyer’s degree of ambiguity aversion. The buyer’s optimization problem is then given by

$$\max_{x \in \Psi} (1 - \varepsilon) \int_0^1 H(c)x(c) \, dc + \varepsilon \inf_{c \in [0,1]} \left\{ x(c)\Delta(c) - \int_c^1 x(u) \, du \right\}, \quad \varepsilon \in (0, 1). \quad (III)$$

Under this preference representation, the buyer faces a trade-off between maximizing her subjective expected utility evaluated at F and limiting the worst-case scenario. As shown above, the former objective is achieved by offering the posted price c^* , whereas the worst-case payoff is maximized by offering the menu characterized in [Proposition 3.1](#). For the remainder of the analysis I assume that $H(c)$ is strictly decreasing in c . Under this assumption, the problem in the absence of ambiguity ($\varepsilon = 0$) is solvable without the use of ironing techniques. Furthermore, I restrict attention to the case where the buyer’s valuation is strictly increasing in the seller’s type. This restriction is less essential but simplifies some of the exposition. The next proposition shows that under these assumptions, the optimal menu is characterized by two parameters: a minimum payoff $\bar{\pi}$ that lies between the minimum payoff under the optimal posted price in the absence of ambiguity, $\bar{\pi}^{\min}$, and the minimum payoff under the menu that perfectly hedges against ambiguity, $\bar{\pi}^{\max}$, as well as a threshold \hat{c} , below which the buyer trades with probability 1 and above which she obtains her minimum payoff or trades with probability 0.

PROPOSITION 3.2. *Assume $H'(c) < 0$ and $v'(c) > 0$ for all $c \in [0, 1]$. Letting x^* denote the solution of problem (III), there exists a threshold $\hat{c} \in [0, c^*]$ and a minimum payoff $\bar{\pi} \in [\bar{\pi}^{\min}, \bar{\pi}^{\max}]$ such that*

- $x^*(c) = 1$ for all $c \in [0, \hat{c}]$,
- $x^*(c)$ is such that $x^*(c)\Delta(c) - \int_0^c x^*(u) \, du = \text{Max}\{\bar{\pi}, 0\}$ for all $c \in (\hat{c}, 1]$.

According to [Proposition 3.2](#), at the optimal mechanism there exists a set of low type sellers, $[0, \hat{c}]$, who are bunched and trade at a base price, while the trading probability of the remaining sellers is such that the buyer’s payoff is constant across $(\hat{c}, 1]$. In particular, if the buyer’s minimum payoff $\bar{\pi}$ associated with the optimal mechanism is weakly negative, types above the threshold \hat{c} trade with probability 0 and the optimal mechanism can be interpreted as a posted price equal to \hat{c} , whereas if the minimum payoff $\bar{\pi}$ is strictly positive, the trading probability for types above the threshold \hat{c} is strictly positive and strictly decreasing. In the latter case, the optimal mechanism separates all types in the interval $(\hat{c}, 1]$. Notice also that if $\hat{c} = c^*$ and $\bar{\pi} = \bar{\pi}_{\min}$, the mechanism characterized in [Proposition 3.2](#) corresponds to the posted price c^* , while if $\hat{c} = 0$ and $\bar{\pi} = \bar{\pi}^{\max}$, it corresponds to the maximally robust menu characterized in [Proposition 3.1](#).

To see why this type of mechanism is optimal for the buyer, suppose that under the posted price c^* , the buyer obtains a positive payoff with all seller types, i.e., $\bar{\pi}^{\min} \geq 0$. This is the case when F is sufficiently skewed to the left and gains from trade with low type sellers are sufficiently large. In such situations, the worst-case scenario for the buyer is when the seller's type is strictly greater than c^* so that the buyer's offer is rejected. Starting from the posted price c^* , suppose now the buyer wishes to increase her minimum payoff. This requires an increase in the trading probability with seller types above c^* , which in turn increases the information rent paid to sellers below c^* . When the information rent paid to these sellers becomes sufficiently large, the buyer's payoff with type $c = 0$ falls below the minimum payoff. If the buyer wishes to increase her minimum payoff further, she thus needs to decrease the trading probability for some types in the bunching region. This is optimally done by reducing the upper bound of that region, i.e., by reducing the threshold \hat{c} . At $\hat{c} = 0$, the buyer's minimum payoff reaches its maximum value, which is equal to $\bar{\pi}^{\max}$.

If, alternatively, $\bar{\pi}^{\min} < 0$ so that the buyer makes losses with some seller types when offering the posted price c^* , increasing the buyer's minimum payoff above $\bar{\pi}^{\min}$ requires directly decreasing the threshold below which the seller trades with probability 1, while keeping the trading probability of the remaining types equal to 0. Thus, so as to increase her minimum payoff, realized when $c = 0$, the buyer initially needs to decrease the posted price to $\hat{c} < c^*$. Once the posted price is such that the buyer's minimum payoff reaches zero, a further increase in the minimum payoff requires an increase in the trading probability for sellers above the threshold \hat{c} , which in turn makes a further decrease of that threshold necessary—again up to the point where $\hat{c} = 0$ and the minimum payoff equals $\bar{\pi}^{\max}$.

The optimal value of \hat{c} and the optimal minimum payoff $\bar{\pi}$ depend on the buyer's confidence parameter ε . The next proposition establishes that the optimal mechanism is a posted price equal to c^* if and only if ε is sufficiently small, while it is the maximally robust menu characterized in Proposition 3.1 if and only if ε is sufficiently large. It also shows that there exists a nonempty region of intermediate values of ε , where the optimal menu combines features of both of these mechanisms.

PROPOSITION 3.3. *Assume $H'(c) < 0$ and $v'(c) > 0$ for all $c \in [0, 1]$. The optimal value of the buyer's minimum payoff $\bar{\pi}$ is increasing in ε , while the optimal value of \hat{c} is decreasing in ε . Moreover, there exist two values $\underline{\varepsilon}$ and $\bar{\varepsilon}$ with $0 \leq \underline{\varepsilon} < \bar{\varepsilon} < 1$ such that*

- $\hat{c} = c^*$ and $\bar{\pi} = \bar{\pi}^{\min}$ if and only if $\varepsilon \leq \underline{\varepsilon}$,
- $\hat{c} = 0$ and $\bar{\pi} = \bar{\pi}^{\max}$ if and only if $\varepsilon \geq \bar{\varepsilon}$.

The result in Proposition 3.3 is very intuitive. The proposition shows that the buyer optimally hedges against ambiguity by separating the seller through the mechanism if and only if she is sufficiently doubtful about the model distribution F . The fact that the threshold of the bunching region, \hat{c} , is decreasing in ε implies that the lower the buyer's confidence in F is, the more separation the optimal mechanism displays.




FIGURE 2. Optimal mechanism: $\bar{\pi}^{\min} < 0$.

The statement of **Proposition 3.3** is shown by using techniques familiar from welfare economics. The buyer maximizes the weighted sum of two payoff functions: the expected payoff $E_F[\pi_b(c)]$ and the minimum payoff $\bar{\pi} = \inf_{c \in [0,1]} \pi_b(c)$. With the structure of the optimal mechanism established in **Proposition 3.2**, it is possible to derive the maximal expected payoff $E_F[\pi_b(c)]$ for each value of $\bar{\pi} \in [\bar{\pi}^{\min}, \bar{\pi}^{\max}]$. The function connecting the two values is strictly decreasing and can be viewed as the analogue of a classic utility-possibility frontier. The optimal mechanism for the buyer is then pinned down by the tangency point between the possibility frontier and the buyer's indifference curve between her expected payoff evaluated at F and the minimum payoff $\bar{\pi}$, which has a slope equal to $-\frac{\epsilon}{1-\epsilon}$.

To see how the optimal mechanism depends on the parameter ϵ , it is useful to distinguish between the two parameter regimes discussed above, where $\bar{\pi}^{\min}$ is positive or negative. Suppose first that $\bar{\pi}^{\min} < 0$: recall that if the buyer's minimum payoff $\bar{\pi}$ lies in the interval $[\bar{\pi}^{\min}, 0]$, the corresponding mechanism is a posted price such that the buyer's payoff with type $c = 0$ equals $\bar{\pi}$, whereas if $\bar{\pi}$ belongs to $(0, \bar{\pi}^{\max}]$, it is a menu that bunches sellers of low type and separates sellers of high type. The proof of **Proposition 3.3** shows that the possibility frontier in this case, illustrated in **Figure 2**, is strictly concave. The slope at the lower bound of its domain, $\bar{\pi}^{\min}$, equals 0, which implies that a posted price equal to c^* is optimal if and only if $\epsilon = 0$. The intuition is that, starting from the posted price c^* , a marginal decrease in the price has no effect on $E_F[\pi_b(c)]$, whereas it has a positive effect on the buyer's minimum payoff. At $\bar{\pi} = 0$, the possibility frontier has a corner. This implies that there exists an interval of values of ϵ for which the optimal value of $\bar{\pi}$ equals 0 and the associated mechanism is a posted price equal to gains from trade at $c = 0$, $\Delta(0)$. This mechanism yields a positive payoff with all sellers in the interval $(0, \Delta(0))$, while it yields a zero payoff with the remaining types. If ϵ exceeds

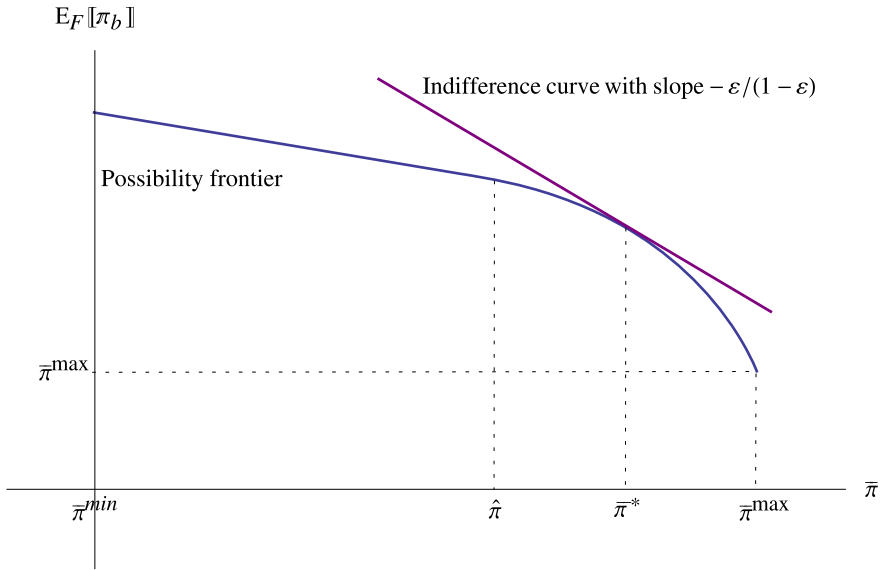


FIGURE 3. Optimal mechanism: $\bar{\pi}^{\min} \geq 0$.

the upper bound of that interval, the optimal value of $\bar{\pi}$ increases continuously in ϵ , up to the point where it reaches $\bar{\pi}^{\max}$ and the optimal mechanism becomes the maximally robust menu characterized in Proposition 3.1. As Proposition 3.3 shows, this value of ϵ is strictly smaller than 1, implying that the maximally robust mechanism is optimal for a nondegenerate set of parameter values.

Considering next the case $\bar{\pi}^{\min} \geq 0$, the utility-possibility frontier has a linear and a strictly concave part as illustrated in Figure 3.² If ϵ is sufficiently small so that the slope of the indifference curve is greater than the slope of the possibility frontier on its linear part, the optimal mechanism is the posted price c^* . At the point where these slopes are the same, the optimal value of $\bar{\pi}$ is indeterminate and then increases continuously in ϵ , up to the point where it reaches $\bar{\pi}^{\max}$, as discussed in the previous case.

The two parameter regimes are equivalent for large values of ϵ , where low types are bunched while high types are separated and the size of the bunching region decreases continuously in ϵ . The main difference arises when ϵ is small. While for $\bar{\pi}^{\min} \geq 0$, the optimal mechanism is a posted price equal to c^* for a nondegenerate set of values of ϵ , for $\bar{\pi}^{\min} < 0$, this is not the case. Instead, the buyer optimally reduces the posted price to a value lower than c^* , even for very small values of ϵ . Appendix B provides a complete characterization of the optimal mechanism for both cases. For the specification with constant gains from trade and a uniform model distribution F , the optimal mechanism can be described in closed form.

PROPOSITION 3.4. Assume $\Delta(c) = \delta$, $\delta < 1$, and $F(c) = c$ for all $c \in [0, 1]$.

²If $c^* > 1$ so that $\bar{\pi}^{\min} > 0$, the linear part disappears. The remaining discussion remains valid.

- *Posted Price.* If $\varepsilon \leq 1 - \frac{1}{\delta}e^{-\frac{1-\delta}{\delta}}$,

$$x^*(c) = \begin{cases} 1 & \text{if } c \leq \delta, \\ 0 & \text{if } c > \delta. \end{cases}$$

- *Partial Separation.* If $\varepsilon \in (1 - \frac{1}{\delta}e^{-\frac{1-\delta}{\delta}}, 1 - \frac{1}{\delta}e^{-\frac{1}{\delta}})$,

$$x^*(c) = \begin{cases} 1 & \text{if } c \leq \hat{c}, \\ \frac{\delta - \hat{c}}{\delta} e^{-\frac{c-\hat{c}}{\delta}} & \text{if } c > \hat{c}, \end{cases}$$

where $\hat{c} = 1 + \delta \ln(\delta(1 - \varepsilon))$.

- *Perfect Separation.* If $\varepsilon \geq 1 - \frac{1}{\delta}e^{-\frac{1}{\delta}}$,

$$x^*(c) = e^{-\frac{c}{\delta}} \quad \forall c.$$

Under the assumption $\Delta(c) = \delta$ and $F(c) = c$, the optimal mechanism in the benchmark case $\varepsilon = 0$ is a posted price equal to δ . The specification thus satisfies $\bar{\pi}^{\min} \geq 0$ and there are three regions of ε to be distinguished. If $\varepsilon \leq \underline{\varepsilon} = 1 - \frac{1}{\delta}e^{-\frac{1-\delta}{\delta}}$, the optimal mechanism is a posted price equal to δ , whereas if $\varepsilon \geq \bar{\varepsilon} = 1 - \frac{1}{\delta}e^{-\frac{1}{\delta}}$, the optimal mechanism is a separating menu that perfectly hedges against ambiguity and yields a safe payoff equal to $\bar{\pi}^{\max} = \delta e^{-\frac{1}{\delta}}$. Alternatively, if ε lies between these two thresholds, the optimal mechanism bunches types below the threshold $\hat{c} = 1 + \delta \ln(\delta(1 - \varepsilon))$ and separates the remaining seller types. The optimal threshold \hat{c} is strictly decreasing in the buyer’s confidence parameter ε , reflecting the hedging function of separating menus in this environment. Figure 4 illustrated the optimal trading probability $x^*(c)$ and the associated payoff $\pi_b(c)$ for different values of ε .

The example helps one to understand why, for intermediate values of ε , the buyer optimally bunches low type sellers and separates high type sellers. Since $\Delta(c) = \delta$ for all $c \in [0, 1]$, the marginal gain of increasing the trading probability with a particular seller, given by the respective gains from trade, is the same across all types. The marginal cost of increasing the trading probability, alternatively, is strictly increasing in the seller’s type because the buyer has to pay the additional information rent to all lower types. As a result, the buyer’s expected payoff evaluated at F is maximized when trading with low type sellers but not with high type sellers. Also the optimal mechanism for a buyer who demands robustness maximizes the trading probability with low type sellers—by letting them trade with probability 1 at a base price—but minimizes the trading probability of high type sellers only up to the point where the minimum payoff is reached.

4. DISCUSSION

4.1 Ambiguity aversion versus risk aversion

Given the longstanding debate on the implications of ambiguity, an important question in this contracting problem is how the effects of ambiguity aversion differ from those of risk aversion. It is well known that if utility functions are not quasilinear, separating




FIGURE 4. Trading probability $x^*(c)$ and payoff $\pi_b(c)$ for different values of $\varepsilon \in [0, 1]$.

menus can be optimal, even when there is no ambiguity. The following discussion shows that the optimality conditions that determine the mechanism under risk aversion are different from those under ambiguity aversion, and provides some intuition for why this is the case.

As an example, suppose the buyer’s utility function u is concave in the difference between her valuation $v(c)$ and the price she pays in exchange for the good. Given the arguments put forward in Section 3, the seller’s payoff as a function of his type and the trading probability $x(c)$ can be written as $\pi_s(c) = \int_c^1 x(t) dt$ so that the price conditional on trading, denoted by $p(c)$, is given by $p(c) = c + \frac{\int_c^1 x(t) dt}{x(c)}$ if $x(c) > 0$ and $p(c) = 0$ otherwise. Normalizing the buyer’s outside option to zero, the risk-averse buyer with single prior belief F thus maximizes

$$\int_0^1 x(c)u\left(\Delta(c) - \int_c^1 \frac{x(t)}{x(c)} dt\right)f(c) dc$$

subject to the monotonicity constraint on $x(c)$. There exist parameter constellations under which the function $x(c)$ that solves this problem corresponds to a separating menu rather than a posted price, as can be verified. However, when the buyer's valuation depends on the seller's type, the cases in which the solution $x(c)$ corresponds exactly to the optimal mechanism under ambiguity aversion, whether in the form of Knightian uncertainty or the ε -contamination model, are nongeneric.

More generally, when the buyer's type increases in the seller's type, the way in which the desire to hedge against ambiguity and to hedge against risk affect the optimal mechanism is typically different. The previous section showed that ambiguity aversion favors the separation of sellers through the mechanism. In particular, the maximally robust mechanism, which is optimal when the buyer is sufficiently ambiguity averse, is such that the trading probability strictly decreases in the seller's type. Such a mechanism still exposes the buyer to risk since the buyer faces a strictly positive probability of not trading. In fact, there does not exist a mechanism that yields a riskless payoff. However, in situations where the buyer's type does not depend too much on the seller's type so that the pooling price yields a positive payoff with all types of seller, the buyer can limit her downside risk by offering the pooling price. If the buyer is sufficiently risk averse, the pooling price is then indeed optimal, implying that risk aversion can favor pooling rather than separation.

As seen in the discussion after [Proposition 3.1](#), the assumption that types are interdependent is crucial for the robust mechanism to be a separating menu. If, alternatively, the buyer's valuation does not depend on the seller's private information, the maximally robust mechanism characterized in [Proposition 3.1](#) is the pooling price. Interestingly, when v is constant, offering the pooling price not only yields an unambiguous but also a riskless payoff for the buyer. Thus, in the independent private value case, hedging against ambiguity can take the same form as hedging against risk.

4.2 Smooth ambiguity aversion

For a given set of measures, the maxmin expected utility model may be seen as a special case of the smooth ambiguity model, developed by [Klibanoff et al. \(2005\)](#), with infinite ambiguity aversion. This section illustrates that the main characteristics of the optimal contract extend to the case of smooth ambiguity aversion and thus do not hinge on the kink property of maxmin expected utility. To keep the analysis tractable, I return to the case of binary types, where the probability distribution of seller types is captured by a single parameter $\sigma \in [0, 1]$. In the smooth ambiguity model, the buyer's utility function is given by

$$E_{\mu}[\Phi(E_{\sigma}[\pi_b])],$$

where $\mu : [0, 1] \rightarrow [0, 1]$ is a subjective prior on a set of probability measures, here captured by $\sigma \in [0, 1]$, and $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ is a function that weighs realizations of the decision maker's expected utility $E_{\sigma}[\pi_b]$. Ambiguity is captured by the second-order belief μ , which measures the buyer's belief about a particular σ being the "correct" probability distribution, while ambiguity attitude is captured by the function Φ . If Φ is linear, the

buyer is ambiguity neutral and her preferences are observationally equivalent to those of a subjective expected utility maximizer. If, alternatively, Φ is concave, the buyer is ambiguity averse and prefers known risks over unknown risks. The degree of ambiguity aversion is measured by the coefficient of absolute ambiguity aversion $-\frac{\Phi'(x)}{\Phi'(x)}$.

As in Section 2, the buyer proposes a menu of the form $\{(1, \alpha c_h + (1 - \alpha)c_l), (\alpha, \alpha c_h)\}$. The optimization problem of the buyer thus amounts to

$$\max_{\alpha \in [0,1]} E_\mu[\Phi(\sigma\alpha(v_h - c_h) + (1 - \sigma)(v_l - \alpha c_h - (1 - \alpha)c_l))].$$

To make ambiguity matter, assume that μ has positive mass on both $[0, \tilde{\sigma})$ and $(\tilde{\sigma}, 1]$, and assume that $\Phi'(\cdot) > 0$ and $\Phi''(\cdot) < 0$. The first-order condition of the buyer's optimization problem is given by

$$\underbrace{E_\mu[\Phi'(\sigma\alpha(v_h - c_h) + (1 - \sigma)(v_l - \alpha c_h - (1 - \alpha)c_l))(\tilde{\sigma} - \sigma) \mid \sigma < \tilde{\sigma}]}_{\equiv \text{MC}(\alpha)} \\ = \underbrace{E_\mu[\Phi'(\sigma\alpha(v_h - c_h) + (1 - \sigma)(v_l - \alpha c_h - (1 - \alpha)c_l))(\sigma - \tilde{\sigma}) \mid \sigma > \tilde{\sigma}]}_{\equiv \text{MG}(\alpha)}.$$

The marginal cost of increasing α , $\text{MC}(\alpha)$, is the marginal decrease in expected utility in the probabilistic scenario that pooling is not optimal ($\sigma < \tilde{\sigma}$), while the marginal gain of increasing α , $\text{MG}(\alpha)$, is the marginal increase in expected utility in the probabilistic scenario that pooling is optimal ($\sigma > \tilde{\sigma}$). Concavity of Φ implies that the marginal cost is increasing in α , whereas the marginal gain is decreasing in α . This implies that there is a unique α that maximizes the buyer's expected payoff. The conditions for an interior solution are

$$\text{MC}(0) < \text{MG}(0) \quad \text{and} \quad \text{MC}(1) > \text{MG}(1).$$

The following proposition summarizes this result.

PROPOSITION 4.1. *The optimal menu of contracts for a buyer with smooth ambiguity aversion is $\{(1, \alpha^* c_h + (1 - \alpha^*)c_l), (\alpha^*, \alpha^* c_h)\}$ with*

$$\alpha^* = \begin{cases} 0 & \text{if } \text{MC}(0) \geq \text{MG}(0), \\ 1 & \text{if } \text{MC}(1) \leq \text{MG}(1), \\ \text{such that } \text{MC}(\alpha^*) = \text{MG}(\alpha^*) & \text{otherwise.} \end{cases}$$

Proposition 4.1 is the smooth counterpart of **Proposition 2.1** in the introductory example. To see the connection to the optimal mechanism under maxmin expected utility, assume that absolute ambiguity aversion is constant, i.e., $-\frac{\Phi'(x)}{\Phi'(x)} = \gamma$. Assume further that the support of μ is $[\underline{\sigma}, \bar{\sigma}]$ and that $\underline{\sigma} < \tilde{\sigma} < \bar{\sigma}$. Under maxmin preferences, the buyer's optimal mechanism, characterized by $\alpha^{\text{MEU}} \equiv \frac{v_l - c_l}{v_h - c_l}$, makes her payoff unambiguous. Under smooth ambiguity aversion, the buyer compromises between maximizing the second-order expectation of her payoff, $E_\mu E_\sigma[\pi_b]$, and limiting her exposure to ambiguity. The expected payoff $E_\mu E_\sigma[\pi_b]$ is maximized by offering a posted price (equal

to either c_l or c_h), whereas ambiguity is eliminated by offering the menu characterized by α^{MEU} . The optimal mechanism under smooth ambiguity aversion is a convex combination of the two. If offering the pooling price maximizes $E_\mu E_\sigma[\pi_b]$, then the optimal contracting parameter α^* lies in the interval $[\alpha^{\text{MEU}}, 1]$; otherwise α^* lies in the interval $[0, \alpha^{\text{MEU}}]$. The more ambiguity averse the buyer is, the closer is α^* to α^{MEU} . This is summarized in Proposition 4.2.

PROPOSITION 4.2. Assume $-\frac{\Phi''(x)}{\Phi'(x)} = \gamma$.

- If $E_\mu[\sigma] < \tilde{\sigma}$, then $\alpha^* \in [0, \alpha^{\text{MEU}}]$ and $\frac{d\alpha^*}{d\gamma} \geq 0$.
- If $E_\mu[\sigma] > \tilde{\sigma}$, then $\alpha^* \in [\alpha^{\text{MEU}}, 1]$ and $\frac{d\alpha^*}{d\gamma} \leq 0$.

5. CONCLUSION

This paper considers the contracting problem between a seller and a buyer, who demands robustness with regard to the distribution of the seller’s private information, in an environment with interdependent values. The analysis shows that if the buyer’s valuation increases in the seller’s valuation, the nature of the optimal mechanism crucially depends on the buyer’s confidence in the underlying type distribution. The maximally robust contract in this situation is a menu that separates all seller types and thereby hedges against the ambiguity perceived by the buyer. As a result, the larger the buyer’s demand for robustness is, the more separation the optimal mechanism displays. This stands in contrast to the case where buyer and seller have independent private values and a posted price is optimal, no matter what the degree of the buyer’s ambiguity aversion is.

One question not addressed in this paper is how the presence of uncertainty in the form of ambiguity affects the efficiency of the equilibrium allocation. In the environment considered, gains from trade are strictly positive with each type of seller and are thus maximized when the buyer offers the pooling price. As the analysis demonstrates, if the buyer knows the underlying type distribution, she offers a posted price potentially smaller than the pooling price, whereas if the buyer perceives and dislikes ambiguity, she proposes a separating menu. How these two mechanisms compare in terms of the social surplus they generate, and, more generally, how the presence of ambiguity affects the welfare of trading parties in environments with asymmetric information might be an interesting question for future research.

APPENDIX A

A.1 Proof of Proposition 2.1

The buyer maximizes $\inf_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} E_\sigma[\pi_b]$. To identify the minimizing prior, consider

$$\frac{\partial E_\sigma[\pi_b]}{\partial \sigma} = \alpha(v_h - c_l) - (v_l - c_l).$$

The buyer's payoff $E_\sigma[\pi_b]$ is decreasing in σ for all $\alpha \leq \frac{v_l - c_l}{v_h - c_l}$ and increasing in σ for all $\alpha \geq \frac{v_l - c_l}{v_h - c_l}$. We thus have $\bar{\sigma} \in \operatorname{arg\,inf}_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} E_\sigma[\pi_b]$ if $\alpha \leq \frac{v_l - c_l}{v_h - c_l}$ holds and $\underline{\sigma} \in \operatorname{arg\,inf}_{\sigma \in [\underline{\sigma}, \bar{\sigma}]} E_\sigma[\pi_b(\alpha)]$ if $\alpha \geq \frac{v_l - c_l}{v_h - c_l}$ holds. The buyer consequently maximizes the step function

$$\Pi = \begin{cases} \bar{\sigma}\alpha(v_h - c_h) + (1 - \bar{\sigma})(v_l - \alpha c_h + (1 - \alpha)c_l) & \text{if } \alpha \leq \frac{v_l - c_l}{v_h - c_l}, \\ \underline{\sigma}\alpha(v_h - c_h) + (1 - \underline{\sigma})(v_l - \alpha c_h + (1 - \alpha)c_l) & \text{if } \alpha > \frac{v_l - c_l}{v_h - c_l}. \end{cases}$$

By definition of $\tilde{\sigma}$, the buyer's payoff $E_\sigma[\pi_b]$ is decreasing in α if $\sigma \leq \tilde{\sigma}$ and increasing in α if $\sigma \geq \tilde{\sigma}$. Consequently, if $\bar{\sigma} \leq \tilde{\sigma}$, both parts of the step function are decreasing in α and Π is maximized at $\alpha = 0$. Similarly, if $\underline{\sigma} \geq \tilde{\sigma}$, both parts of the step function are increasing in α and Π is maximized at $\alpha = 1$. If $\underline{\sigma} < \tilde{\sigma} < \bar{\sigma}$, Π is strictly increasing in α on the interval $[0, \frac{v_l - c_l}{v_h - c_l}]$ and strictly decreasing in α on the interval $[\frac{v_l - c_l}{v_h - c_l}, 1]$, and therefore is maximized at $\frac{v_l - c_l}{v_h - c_l}$.

A.2 Proof of Proposition 3.1

Suppose that the solution of problem (II) is such that $\pi_b(c)$ is constant across c . This implies that $x(c)$ is differentiable in c . Indeed, notice that, given $\pi_b(c) = C$ for some $C \in \mathbb{R}$, we have $x(c)\Delta(c) = \pi_s(c) + C$ for all $c \in [0, 1]$. As π_s is absolutely continuous, a standard consequence of the seller's incentive constraints, $x(c)$ is continuous in c . But as $\pi_s(c) = \int_c^1 x(u) du$ and $\Delta(c)$ is differentiable and bounded away from zero, we get from the fundamental theorem of calculus that $x(c)$ is differentiable in c .

We then have $\pi'_b(c) = x'(c)\Delta(c) + x(c)v'(c)$, so the optimal mechanism solves the differential equation

$$x'(c)\Delta(c) + x(c)v'(c) = 0. \tag{1}$$

The solution to (1) is given by $x(c) = D \exp[-\int_0^c \frac{v'(t)}{\Delta(t)} dt]$, $D \in \mathbb{R}$. Together with the condition $x(0) = 1$, this yields $D = 1$ and hence $x^*(c)$ as characterized in Proposition 3.1. Let the associated (constant) payoff be denoted by $\bar{\pi}^{\max}$.

We then need to show that any optimal mechanism indeed equalizes the buyer's payoff across c . Consider a nonincreasing function $\tilde{x} : [0, 1] \rightarrow [0, 1]$ that satisfies $\tilde{x}(c)\Delta(c) - \int_c^1 \tilde{x}(u) du \geq \bar{\pi}^{\max}$. We can first show that $\tilde{x}(c) \geq x^*(c)$ for all c . For each $c \in [0, 1]$, let $\tilde{x}_c : [0, 1] \rightarrow \mathbb{R}$ be an auxiliary function, defined by $\tilde{x}_c(u) = D_c \exp[-\int_0^u \frac{v'(t)}{\Delta(t)} dt]$ with D_c such that $\tilde{x}_c(c) = \tilde{x}(c)$. Since $x^*(u) = \exp[-\int_0^u \frac{v'(t)}{\Delta(t)} dt]$, for all c such that $\tilde{x}(c) < x^*(c)$, we have $D_c < 1$ and hence $\tilde{x}_c(u) < x^*(u) \forall u \in [0, 1]$.

Toward a contradiction, suppose now there exists some $c \in [0, 1]$ such that $\tilde{x}(c) < x^*(c)$. Notice that since the payoff with the highest type, $\tilde{x}(1)\Delta(1)$, must be weakly greater than $\bar{\pi}^{\max} = x^*(1)\Delta(1)$, we have $\tilde{x}(1) \geq x^*(1)$ and hence $D_1 \geq 1$. Now we know that $D_c < 1$ since by assumption $\tilde{x}(c) < x^*(c)$. We therefore have $\tilde{x}(1) > \tilde{x}_c(1)$. Since \tilde{x} is nonincreasing and \tilde{x}_c is continuous, we then have $\tilde{x}(u) > \tilde{x}_c(u)$ on a left neighborhood of 1. Let c' be the supremum of the set $\{u \in [0, 1] : \tilde{x}(u) \leq \tilde{x}_c(u)\}$, which is well defined

as $\tilde{x}(c) = \tilde{x}_c(c)$. Since \tilde{x} is nonincreasing, it must then hold that $\tilde{x}(c') \leq \tilde{x}_c(c') < x^*(c')$. Considering the buyer's payoff with type c' , we obtain

$$\tilde{x}(c') - \int_{c'}^1 \tilde{x}(u) \, du \leq \tilde{x}_{c'}(c') - \int_{c'}^1 \tilde{x}_{c'}(u) \, du = \tilde{x}_{c'}(1)\Delta(1) < x^*(1)\Delta(1) = \bar{\pi}^{\max}.$$

The first inequality follows from the facts that $\tilde{x}(c') = \tilde{x}_{c'}(c')$ (by definition of \tilde{x}_c) and $\tilde{x}(u) > \tilde{x}_{c'}(u)$ for all $u > c'$. The following equality follows comes from the property that the buyer's payoff is constant in c under $\tilde{x}_{c'}$. The last inequality follows from $\tilde{x}(c') < x^*(c')$, which implies $D_{c'} < 1$ and hence $\tilde{x}_{c'}(u) < x^*(u)$ for all $u \in [0, 1]$. Together, this contradicts our initial assumption $\tilde{x}(c)\Delta(c) - \int_c^1 \tilde{x}(u) \, du \geq \bar{\pi}^{\max}$ for all c and therefore implies $\tilde{x}(c) \geq x^*(c)$ for all c .

Finally, consider the buyer's payoff with type $c = 0$. Given that the minimum payoff associated to \tilde{x} must be weakly greater than $\bar{\pi}^{\max}$, it must hold that

$$\tilde{x}(0)\Delta(0) - \int_0^1 \tilde{x}(u) \, du \geq \Delta(0) - \int_0^1 x^*(u) \, du = \bar{\pi}^{\max}.$$

Since $\tilde{x}(0)\Delta(0) \leq \Delta(0)$ and $\tilde{x}(c) \geq x^*(c)$ for all c , this inequality can only be satisfied if $\tilde{x}(c) = x^*(c)$ for all c .

A.3 Proof of Proposition 3.2

To characterize the solution of problem (III), consider the auxiliary optimization problem, where the endogenous minimum payoff is treated as an exogenous parameter $\bar{\pi} \in \mathbb{R}$:

$$\max_{x \in \Psi} \int_0^1 H(c)x(c) \, dc \tag{III'}$$

subject to the pointwise constraints

$$x(c)\Delta(c) - \int_c^1 x(u) \, du \geq \bar{\pi} \quad \forall c \in [0, 1]. \tag{2}$$

Evidently, any solution of problem (III) must also be a solution of (III') for some $\bar{\pi}$, as otherwise the buyer could increase his expected payoff evaluated at F while maintaining the same minimum payoff $\inf_{c \in [0, 1]} \{x(c)\Delta(c) - \int_c^1 x(u) \, du\}$. We can restrict our attention to values of $\bar{\pi}$ that are weakly greater than the minimum payoff under the mechanism that maximizes $\int_0^1 H(c)x(c) \, dc$, i.e., that solves problem (I). Recall that this payoff is given by $\bar{\pi}^{\min} = \max\{\Delta(0) - c^*, 0\}$. Moreover, the minimum payoff $\bar{\pi}$ has to be weakly smaller than the maximum value of $\inf_{c \in [0, 1]} \{x(c)\Delta(c) - \int_c^1 x(u) \, du\}$, attained at the mechanism that solves problem (II), as otherwise the feasible set is empty. This payoff is given by $\bar{\pi}^{\max} = \Delta(1) \exp(-\int_0^1 \frac{v'(t)}{\Delta(t)} \, dt)$.

It will be useful to derive the function $x_{\bar{\pi}}^{\min} : [0, 1] \rightarrow [0, 1]$ under which the buyer's payoff $\pi_b(c)$ is equal to $\bar{\pi}$ if $\bar{\pi} > 0$ and equal to 0 if $\bar{\pi} \leq 0$ for all $c \in [0, 1]$. This requires that the buyer's payoff is constant across $c \in [0, 1]$ and therefore that $x_{\bar{\pi}}^{\min}(c) = D \exp[-\int_0^c \frac{v'(t)}{\Delta(t)} \, dt]$ for some $D \in \mathbb{R}$ (see Appendix A.2). The associated (constant) payoff

is equal to $\max\{\bar{\pi}, 0\}$ if $x_{\bar{\pi}}^{\min}(1)\Delta(1) = \max\{\bar{\pi}, 0\}$, i.e., if $D = \frac{\max\{\bar{\pi}, 0\}}{\Delta(1)} \exp(\int_0^1 \frac{v'(t)}{\Delta(t)} dt)$. We thus have

$$x_{\bar{\pi}}^{\min}(c) = \frac{\max\{\bar{\pi}, 0\}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right).$$

We can first show that the pointwise constraints (2) are satisfied only if $x(c) \geq x_{\bar{\pi}}^{\min}(c)$ for all $c \in [0, 1]$. First, if $\bar{\pi} \leq 0$ so that $x_{\bar{\pi}}^{\min}(c) = 0$ for all $c \in [0, 1]$, $x(c) \geq x_{\bar{\pi}}^{\min}(c)$ must be trivially satisfied. Second, if $\bar{\pi} > 0$ and $x(c) < x_{\bar{\pi}}^{\min}(c)$ for some $c \in [0, 1]$, by an analogous argument to that in [Appendix A.2](#), there exists some c' and some function $x_{c'}(u) = D_{c'} \exp[-\int_0^u \frac{v'(t)}{\Delta(t)} dt]$ with $D_{c'}$ such that $x_{c'}(c) = x(c)$ such that

$$x(c') - \int_{c'}^1 x(u) du \leq x_{c'}(c') - \int_{c'}^1 x_{c'}(u) du = x_{c'}(1)\Delta(1) < x_{\bar{\pi}}^{\min}(1)\Delta(1) = \bar{\pi},$$

thus violating (2).

Suppose now a solution to (III') exists (which we are going to show later) and let it be denoted by $x_{\bar{\pi}}$. We can then demonstrate that for each $c \in [0, 1]$, either $x_{\bar{\pi}}(c) = 1$ or $x_{\bar{\pi}}(c) = x_{\bar{\pi}}^{\min}(c)$. To see this, let $C = \{c \in [0, 1] : x_{\bar{\pi}}(c) \in (x_{\bar{\pi}}^{\min}(c), 1)\}$ and consider the function $\hat{x} : [0, 1] \rightarrow [0, 1]$, defined by

$$\hat{x}(c) = \begin{cases} 1 & \text{if } c \leq \hat{c}, \\ x_{\bar{\pi}}^{\min}(c) & \text{if } c > \hat{c}, \end{cases}$$

where $\hat{c} \in (0, 1)$ is such that $\int_0^1 \hat{x}(u) du = \int_0^1 x_{\bar{\pi}}(u) du$. We can first verify that \hat{x} satisfies the pointwise constraints (2). For $c \leq \hat{c}$, where $\hat{x}(c) = 1$, notice that the buyer's payoff, given by $\Delta(c) - \int_c^1 \hat{x}(u) du = \Delta(c) - (\hat{c} - c) - \int_{\hat{c}}^1 x_{\bar{\pi}}^{\min}(u) du$, strictly increases in c , so that the only relevant constraint is $\Delta(0) - \int_0^1 \hat{x}(u) du \geq \bar{\pi}$. Since $\int_0^1 \hat{x}(u) du = \int_0^1 x_{\bar{\pi}}(u) du$, this constraint is indeed satisfied. For $c > \hat{c}$, constraints (2) are satisfied by construction.

Define next the functions $\hat{X}(c) \equiv \int_0^c \hat{x}(u) du$ and $X_{\bar{\pi}}(c) \equiv \int_0^c x_{\bar{\pi}}(u) du$, and consider their difference $\hat{X}(c) - X_{\bar{\pi}}(c)$. Since $\hat{x}(c) \geq x_{\bar{\pi}}(c)$ for all $c \leq \hat{c}$, $\hat{X}(c) - X_{\bar{\pi}}(c)$ is increasing on $[0, \hat{c}]$, strictly so on the interior of C .³ Similarly, since $\hat{x}(c) \leq x_{\bar{\pi}}(c)$ for all $c \geq \hat{c}$, $\hat{X}(c) - X_{\bar{\pi}}(c)$ is decreasing on $[\hat{c}, 1]$, again strictly on the interior of C . This together with $\hat{X}(0) = X_{\bar{\pi}}(0)$ and $\hat{X}(1) = X_{\bar{\pi}}(1)$ implies that $\hat{X}(c) - X_{\bar{\pi}}(c) \geq 0$ for all $c \in [0, 1]$ and $\hat{X}(c) - X_{\bar{\pi}}(c) > 0$ for all $c \in \text{int } C$.

Consider then the difference in the buyer's expected payoff evaluated at F associated to \hat{x} and $x_{\bar{\pi}}$. After an integration by parts, we obtain

$$\begin{aligned} & \int_0^1 H(c)\hat{x}(c) dc - \int_0^1 H(c)x_{\bar{\pi}}(c) dc \\ &= [H(c)[\hat{X}(c) - X_{\bar{\pi}}(c)]]_0^1 - \int_0^1 H'(c)[\hat{X}(c) - X_{\bar{\pi}}(c)] dc. \end{aligned}$$

The first term on the right-hand side equals 0. Given $H'(c) < 0$ and $\hat{X}(c) \geq X_{\bar{\pi}}(c)$ for all $c \in [0, 1]$ and $\hat{X}(c) > X_{\bar{\pi}}(c)$ for all $c \in \text{int } C$, the term $\int_0^1 H'(c)(\hat{X}(c) - X_{\bar{\pi}}(c)) dc$

³Notice that $\frac{d\hat{X}(c) - X_{\bar{\pi}}(c)}{dc} = \hat{x}(c) - x_{\bar{\pi}}(c)$ almost everywhere.

is nonnegative, and hence $\int_0^1 H(c)\hat{x}(c) dc \leq \int_0^1 H(c)x_{\bar{\pi}}(c) dc$, only if the set $\text{int } C$ is empty.

The previous argument, together with the monotonicity constraint on $x_{\bar{\pi}}$ implies that there exists a threshold \hat{c} such that $x_{\bar{\pi}}(c) = 1$ for all $c \leq \hat{c}$ and $x_{\bar{\pi}}(c) = x_{\bar{\pi}}^{\min}(c)$ for all $c > \hat{c}$. Notice that, in principle, the function $x_{\bar{\pi}}$ can take infinitely many values at \hat{c} . This is the only degree of flexibility, so setting $x = 1$ at \hat{c} is without loss of generality. The threshold \hat{c} maximizes the buyer's expected payoff evaluated at F ,

$$\int_0^{\hat{c}} H(c) dc + \int_{\hat{c}}^1 H(c)x_{\bar{\pi}}^{\min}(c) dc, \tag{3}$$

subject to the pointwise constraints (2). Given the structure of $x_{\bar{\pi}}$, the only relevant of those constraints is that for $c = 0$, as argued above. The first derivative of (3) with respect to \hat{c} , given by $H(\hat{c})(1 - x_{\bar{\pi}}^{\min}(\hat{c}))$, shows that (3) is strictly increasing in \hat{c} on $[0, c^*]$ and strictly decreasing in \hat{c} on $[c^*, 0]$. The optimal mechanism is therefore characterized by the maximal value of \hat{c} on $[0, c^*]$ such that the buyer's payoff with type $c = 0$, given by $\Delta(0) - \hat{c} - \int_{\hat{c}}^1 x_{\bar{\pi}}^{\min}(u) du$, is weakly greater than $\bar{\pi}$. If this condition is satisfied at $\hat{c} = c^*$, i.e., if

$$\Delta(0) - c^* - \int_{c^*}^1 \frac{\max\{0, \bar{\pi}\}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc \geq \bar{\pi}, \tag{4}$$

the optimal mechanism is characterized by the threshold c^* . Otherwise the optimal threshold is (uniquely) determined by the condition

$$\Delta(0) - \hat{c} - \int_{\hat{c}}^1 \frac{\max\{0, \bar{\pi}\}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc = \bar{\pi}. \tag{5}$$

Problem (III') is thus uniquely solved by the function

$$x_{\bar{\pi}}(c) = \begin{cases} 1 & \text{if } c \leq \hat{c}, \\ \frac{\max\{0, \bar{\pi}\}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) & \text{if } c > \hat{c}, \end{cases} \tag{6}$$

where $\hat{c} = c^*$ if (4) is satisfied and \hat{c} is such that (5) holds otherwise.

A.4 Proof of Proposition 3.3

Since a solution of problem (III) must also be a solution of (III') for some $\bar{\pi} \in [\bar{\pi}^{\min}, \bar{\pi}^{\max}]$, to solve the problem (III), we can consider the simpler problem where the buyer chooses from the collection of functions $x_{\bar{\pi}}$ as defined in (6) with $\bar{\pi} \in [\bar{\pi}^{\min}, \bar{\pi}^{\max}]$. The buyer's optimization problem can thus be written as

$$\max_{x_{\bar{\pi}}} (1 - \varepsilon) \int_0^1 H(c)x_{\bar{\pi}}(c) dc + \varepsilon \inf_{c \in [0,1]} \left\{ x_{\bar{\pi}}(c)\Delta(c) - \int_c^1 x_{\bar{\pi}}(u) du \right\}. \tag{7}$$

Define $\Pi_F(\bar{\pi}) = \int_0^1 H(c)x_{\bar{\pi}}(c) dc$ and notice that $\inf_{c \in [0,1]} \{x_{\bar{\pi}}(c)\Delta(c) - \int_c^1 x_{\bar{\pi}}(u) du\} = \bar{\pi}$. Problem (7) is then analogous to the problem of maximizing $\varepsilon\Pi_F + (1 - \varepsilon)\bar{\pi}$ subject to

$\Pi_F = \Pi_F(\bar{\pi})$, where $\Pi_F(\bar{\pi})$ is the maximal expected payoff evaluated at F as a function of the minimum payoff $\bar{\pi}$, analogous to a conventional utility-possibility frontier. The optimal mechanism is then determined by the tangency points of the buyer's indifference curves of $\varepsilon\Pi_F + (1 - \varepsilon)\bar{\pi}$ and the possibility constraint $\Pi_F(\bar{\pi})$. The indifference curves are straight lines with a slope equal to $-\frac{\varepsilon}{1-\varepsilon}$.

So as to derive the tangency points, it will be useful to distinguish the two cases where $\bar{\pi}^{\min}$ is positive or negative. Suppose first $\bar{\pi}^{\min} \geq 0$ and let $\hat{\pi}$ denote the value of $\bar{\pi}$ such that (4) is satisfied with equality. Below $\hat{\pi}$ the threshold \hat{c} is equal to c^* , while above $\hat{\pi}$ the value of \hat{c} is determined by (5). Since this condition depends on $\bar{\pi}$, to analyze $\Pi_F(\bar{\pi})$, we need to find the derivative of \hat{c} with respect to $\bar{\pi}$. Taking the total differential of (5), we obtain

$$\frac{d\hat{c}}{d\bar{\pi}} = -\frac{1 + \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc}{1 - \frac{\bar{\pi}}{\Delta(1)} \exp\left(\int_{\hat{c}}^1 \frac{v'(t)}{\Delta(t)} dt\right)}.$$

With this, we have

$$\begin{aligned} \Pi'_F(\bar{\pi}) &= \begin{cases} \frac{1}{\Delta(1)} \int_{c^*}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) H(c) dc & \text{if } \bar{\pi} \in [\bar{\pi}^{\min}, \hat{\pi}], \\ -H(\hat{c}) - \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) (H(\hat{c}) - H(c)) dc & \text{if } \bar{\pi} \in (\hat{\pi}, \bar{\pi}^{\max}], \end{cases} \quad (8) \\ \Pi''_F(\bar{\pi}) &= \begin{cases} 0 & \text{if } \bar{\pi} \in [\bar{\pi}^{\min}, \hat{\pi}], \\ -H'(\hat{c}) \left(1 + \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc\right) \frac{d\hat{c}}{d\bar{\pi}} & \text{if } \bar{\pi} \in (\hat{\pi}, \bar{\pi}^{\max}]. \end{cases} \end{aligned}$$

It can be verified that the derivative of $\Pi_F(\bar{\pi})$ at $\bar{\pi} = \hat{\pi}$ exists and that $\Pi'_F(\bar{\pi})$ is strictly negative and weakly decreasing in $\bar{\pi}$. The function $\Pi_F(\bar{\pi})$ is thus differentiable, strictly decreasing, and weakly concave on its domain. More specifically, $\Pi_F(\bar{\pi})$ is linear on $[0, \hat{\pi}]$ and strictly concave on $[\hat{\pi}, \bar{\pi}^{\max}]$. This implies that there exists only one parameter point at which the buyer's problem has multiple solutions, namely when ε is such that $-\frac{\varepsilon}{1-\varepsilon} = \Pi'_F(\bar{\pi})$, $\bar{\pi} \in [0, \hat{\pi}]$.

We can then derive the thresholds $\underline{\varepsilon}$ and $\bar{\varepsilon}$. The optimal mechanism is characterized by $\bar{\pi} = 0$ if $\Pi'_F(0) \leq -\frac{\varepsilon}{1-\varepsilon}$, while it is characterized by $\bar{\pi} = \bar{\pi}^{\max}$ if $\Pi'_F(\bar{\pi}^{\max}) \geq -\frac{\varepsilon}{1-\varepsilon}$. We thus have $\underline{\varepsilon} = \frac{-\Pi'_F(0)}{1-\Pi'_F(0)}$ and $\bar{\varepsilon} = \frac{-\Pi'_F(\bar{\pi}^{\max})}{1-\Pi'_F(\bar{\pi}^{\max})}$. Since $\Pi'_F(0)$ and $\Pi'_F(\bar{\pi}^{\max})$ take finite, nonzero values, we have $0 < \underline{\varepsilon} < \bar{\varepsilon} < 1$.

Consider next the case $\bar{\pi}^{\min} < 0$. Under this specification, $\hat{\pi} = \bar{\pi}^{\min}$. The threshold \hat{c} is therefore always determined by condition (5). For weakly negative values of $\bar{\pi}$, $x_{\bar{\pi}}(c)$ is a function with a single step at $\Delta(0) - \hat{c}$ and so $\Pi_F(\bar{\pi}) = \int_0^{\Delta(0) - \bar{\pi}} H(c) dc$ for all $\bar{\pi} \leq 0$. The first and second derivatives of $\Pi_F(\bar{\pi})$ in this case are

$$\Pi'_F(\bar{\pi}) = -H(\Delta(0) - \bar{\pi}) < 0 \quad \text{and} \quad \Pi''_F(\bar{\pi}) = H'(\Delta(0) - \bar{\pi}) < 0.$$

On $[\bar{\pi}^{\min}, 0]$, $\Pi_F(\bar{\pi})$ is thus strictly decreasing and strictly concave. The properties of $\Pi_F(\bar{\pi})$ for $\bar{\pi} > 0$ have been described above. At $\bar{\pi} = 0$, the function $\Pi_F(\bar{\pi})$ has a corner.

To see this, notice that $\Pi'_{F(-)}(0) = -H(\Delta(0))$ and

$$\begin{aligned} \Pi'_{F(+)}(0) &= -H(\Delta(0)) - \frac{1}{\Delta(1)} \int_{\Delta(0)}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) \underbrace{(H(\Delta(0)) - H(c))}_{>0} dc \\ &< -H(\Delta(0)). \end{aligned}$$

Since $\Pi'_{F(-)}(0) > \Pi'_{F(+)}(0)$, the derivative function of $\Pi_F(\bar{\pi})$ is strictly decreasing in $\bar{\pi}$. The function $\Pi_F(\bar{\pi})$ is thus strictly concave and the solution of the buyer's optimization problem is consequently unique.

Again we can derive the thresholds $\underline{\varepsilon}$ and $\bar{\varepsilon}$. Given that $\Pi'_F(\bar{\pi}^{\min}) = -H(c^*) = 0$, the condition $\Pi'_F(\bar{\pi}^{\min}) \geq -\frac{\varepsilon}{1-\varepsilon}$ can only be satisfied when $\varepsilon = 0$. We thus have $\underline{\varepsilon} = 0$ and, equivalently to the previous case, $\bar{\varepsilon} = \frac{-\Pi'_F(\bar{\pi}^{\max})}{1-\Pi'_F(\bar{\pi}^{\max})} < 1$.

A.5 Proof of Proposition 3.4

For the specification $\Delta(c) = \delta$, $\delta < 1$, and $F(c) = c$, we have $H(c) = \delta - c$ and hence $c^* = \delta$. As can be verified, this implies $\bar{\pi}^{\min} = \hat{\pi} = 0$. Thus, the threshold \hat{c} is determined by condition (5). Under the stated assumptions, this condition simplifies to $\bar{\pi} = (\delta - \hat{c})e^{-\frac{1-\hat{c}}{\delta}}$. We then obtain

$$x_{\bar{\pi}}(c) = \begin{cases} 1 & \text{if } c \leq \hat{c}, \\ \frac{1}{\delta} \bar{\pi} \exp\left(\frac{1-c}{\delta}\right) & \text{if } c > \hat{c} \end{cases}$$

and thus

$$\Pi_F(\bar{\pi}) = \int_0^{\hat{c}} H(c) dc + \int_{\hat{c}}^1 \frac{1}{\delta} \bar{\pi} \exp\left(\frac{1-c}{\delta}\right) H(c) dc, \quad \Pi'_F(\bar{\pi}) = 1 - \delta \exp\left(\frac{1-\hat{c}}{\delta}\right),$$

with \hat{c} such that $\bar{\pi} = (\delta - \hat{c})e^{-\frac{1-\hat{c}}{\delta}}$ for all $\bar{\pi} \in [\bar{\pi}^{\min}, \bar{\pi}^{\max}]$.

We can then derive the thresholds $\underline{\varepsilon} = \frac{-\Pi'_F(0)}{1-\Pi'_F(0)}$ and $\bar{\varepsilon} = \frac{-\Pi'_F(\bar{\pi}^{\max})}{1-\Pi'_F(\bar{\pi}^{\max})}$. Noticing that according to (5), $\hat{c} = \delta$ at $\bar{\pi} = 0$ and $\hat{c} = 0$ at $\bar{\pi} = \bar{\pi}^{\max}$, we have

$$\underline{\varepsilon} = 1 - \frac{1}{\delta} \exp\left(-\frac{1-\delta}{\delta}\right), \quad \bar{\varepsilon} = 1 - \frac{1}{\delta} \exp\left(-\frac{1}{\delta}\right).$$

With $\bar{\pi}^{\max} = \delta e^{-\frac{1}{\delta}}$, the optimal mechanism for $\varepsilon \leq \underline{\varepsilon}$ and $\varepsilon \geq \bar{\varepsilon}$ is then, respectively, characterized by

$$x^*(c) = \begin{cases} 1 & \text{if } c \leq \delta, \\ 0 & \text{if } c > \delta \end{cases} \quad \text{and} \quad x^*(c) = e^{-\frac{c}{\delta}}, \quad c \in [0, 1].$$

Finally, we consider the case $\varepsilon \in (\underline{\varepsilon}, \bar{\varepsilon})$. Here the optimal mechanism is characterized by the condition $-\Pi'_F(\bar{\pi}) = \frac{\varepsilon}{1-\varepsilon}$; that is,

$$-1 + \delta \exp\left(\frac{1-\hat{c}}{\delta}\right) = \frac{\varepsilon}{1-\varepsilon}.$$

Solving for \hat{c} yields $\hat{c} = 1 + \delta \ln(\delta(1 - \varepsilon))$. Recalling that $\bar{\pi} = (\delta - \hat{c})e^{-\frac{1-\hat{c}}{\delta}}$, the optimal mechanism is characterized by

$$x^*(c) = \begin{cases} 1 & \text{if } c \leq \hat{c}, \\ \frac{\delta - \hat{c}}{\delta} e^{-\frac{c-\hat{c}}{\delta}} & \text{if } c > \hat{c}, \end{cases}$$

where $\hat{c} = 1 + \delta \ln(\delta(1 - \varepsilon))$.

A.6 Proof of Proposition 4.2

Under constant absolute ambiguity aversion, we have $\Phi(x) = -\frac{1}{\gamma}e^{-\gamma x}$ (see [Klibanoff et al. 2005](#)).

Suppose first that $E_\mu[\sigma] > \tilde{\sigma}$ holds. Note that $E_\mu E_\sigma[\pi_b]$ increases in α and that $\text{Var}_\mu[E_\sigma[\pi_b]] > 0$ for all $\alpha < \alpha^{\text{MEU}}$. This implies that the expected utility distribution induced by any $\alpha < \alpha^{\text{MEU}}$ is second-order stochastically dominated by the (degenerate) distribution induced by α^{MEU} . Since Φ is strictly concave, this implies $\alpha^* \in [\alpha^{\text{MEU}}, 1]$.

Letting $\Pi_\sigma(\alpha) = \sigma\alpha(v_h - c_h) + (1 - \sigma)(v_l - \alpha c_h - (1 - \alpha)c_l)$, the optimal value of α^* is characterized by

$$E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma})] = 0 \tag{9}$$

if the solution is interior (otherwise a marginal change in γ has no effect). Taking the total differential yields

$$\frac{d\alpha^*}{d\gamma} = -\frac{E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma})\Pi_\sigma(\alpha^*)]}{\gamma E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma})^2]},$$

$\frac{d\alpha^*}{d\gamma} \leq 0$ if the numerator is weakly positive on the interval $[\alpha^{\text{MEU}}, 1]$. Suppose not. Then

$$\begin{aligned} & E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\tilde{\sigma} - \sigma)\Pi_\sigma(\alpha^*) \mid \sigma \leq \tilde{\sigma}] \\ & > E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma})\Pi_\sigma(\alpha^*) \mid \sigma \geq \tilde{\sigma}]. \end{aligned}$$

But

$$\begin{aligned} & E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\tilde{\sigma} - \sigma)\Pi_\sigma(\alpha^*) \mid \sigma \leq \tilde{\sigma}] \\ & \leq \Pi_{\tilde{\sigma}}(\alpha^*) E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\tilde{\sigma} - \sigma) \mid \sigma \leq \tilde{\sigma}] \\ & = \Pi_{\tilde{\sigma}}(\alpha^*) E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma}) \mid \sigma \geq \tilde{\sigma}] \\ & \leq E_\mu[\exp(-\gamma\Pi_\sigma(\alpha^*))(\sigma - \tilde{\sigma})\Pi_\sigma(\alpha^*) \mid \sigma \geq \tilde{\sigma}], \end{aligned}$$

which follows from (9) and the fact that $\frac{\partial \Pi_\sigma(\alpha)}{\partial \sigma} > 0 \forall \alpha \in [\alpha^{\text{MEU}}, 1]$ —a contradiction. Hence, $\frac{d\alpha^*}{d\gamma} \leq 0$.

The proof for the case $E_\mu[\sigma] < \tilde{\sigma}$ is analogous, where the above inequalities are reversed. Just note that second-order stochastic dominance implies $\alpha^* \in [0, \alpha^{\text{MEU}}]$ and that $\frac{\partial \Pi_\sigma(\alpha)}{\partial \sigma} < 0 \forall \alpha \in [0, \alpha^{\text{MEU}}]$.

APPENDIX B: FULL CHARACTERIZATION OF THE OPTIMAL MECHANISM

PROPOSITION B.1. Assume $H'(c) < 0$ and $v'(c) > 0$ for all $c \in [0, 1]$. If $\bar{\pi}^{\min} \geq 0$, the solution of problem (III) is generically unique and the optimal mechanism is characterized by

- $\bar{\pi} = 0$ and $\hat{c} = c^*$ if $\varepsilon < \underline{\varepsilon}$,
- $\bar{\pi} \in [0, \hat{\pi}]$ and $\hat{c} = c^*$ if $\varepsilon = \underline{\varepsilon}$, where $\hat{\pi}$ is such that

$$\hat{\pi} = \Delta(0) - c^* - \int_{c^*}^1 \frac{\hat{\pi}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc,$$

- $\bar{\pi}$ and \hat{c} such that

$$\frac{\varepsilon}{1 - \varepsilon} = H(\hat{c}) + \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) (H(\hat{c}) - H(c)) dc,$$

$$\bar{\pi} = \Delta(0) - \hat{c} - \int_{\hat{c}}^1 \frac{\bar{\pi}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) dc$$

if $\varepsilon \in (\underline{\varepsilon}, \bar{\varepsilon})$,

- $\bar{\pi} = \bar{\pi}^{\max}$ and $\hat{c} = 0$ if $\varepsilon \geq \bar{\varepsilon}$,

where $\underline{\varepsilon}$ and $\bar{\varepsilon}$, respectively, are defined by

$$\frac{\underline{\varepsilon}}{1 - \underline{\varepsilon}} = -\frac{1}{\Delta(1)} \int_{c^*}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) H(c) dc,$$

$$\frac{\bar{\varepsilon}}{1 - \bar{\varepsilon}} = H(0) + \frac{1}{\Delta(1)} \int_0^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) (H(0) - H(c)) dc.$$

PROOF. The derivative function of $\Pi_F(\bar{\pi})$ is as derived in (8). The threshold $\underline{\varepsilon}$ is such that the slope of the indifference curve coincides with the derivative of the possibility frontier at $\bar{\pi} \in [\bar{\pi}^{\min}, \hat{\pi}]$:

$$-\frac{\underline{\varepsilon}}{1 - \underline{\varepsilon}} = \frac{1}{\Delta(1)} \int_{c^*}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) H(c) dc.$$

For all $\varepsilon \leq \underline{\varepsilon}$, we then have $\hat{c} = c^*$ and $\bar{\pi} = 0$.

The threshold $\bar{\varepsilon}$ is such that the slope of the indifference curve coincides with the derivative of the possibility frontier at $\bar{\pi} = \bar{\pi}^{\max}$. At $\bar{\pi} = \bar{\pi}^{\max}$, condition (5) yields $\hat{c} = 0$. Given this, the threshold $\bar{\varepsilon}$ is characterized by

$$-\frac{\bar{\varepsilon}}{1 - \bar{\varepsilon}} = -H(0) - \frac{1}{\Delta(1)} \int_0^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)} dt\right) (H(0) - H(c)) dc.$$

For $\varepsilon \geq \bar{\varepsilon}$, we therefore have $\hat{c} = 0$ and $\bar{\pi} = \bar{\pi}^{\max}$.

Finally, if $\varepsilon \in (\underline{\varepsilon}, \bar{\varepsilon})$, the optimal value of $\bar{\pi}$ is pinned down by the tangency point between the indifference curve and the possibility frontier on its strictly concave part,

$$-\frac{\varepsilon}{1-\varepsilon} = -H(\hat{c}) - \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) (H(\hat{c}) - H(c)) dc,$$

where \hat{c} is determined by (5). □

PROPOSITION B.2. *Assume $H'(c) < 0$ and $v'(c) > 0$ for all $c \in [0, 1]$. If $\bar{\pi}^{\min} < 0$, the solution of problem (III) is unique and the optimal mechanism is characterized by*

- $\bar{\pi} = \Delta(0) - \hat{c}$ with \hat{c} such that $H(\hat{c}) = \frac{\varepsilon}{1-\varepsilon}$ if $\varepsilon < \varepsilon'$,
- $\bar{\pi} = 0$ and $\hat{c} = \Delta(0)$ if $\varepsilon \in [\varepsilon', \varepsilon'']$,
- $\bar{\pi}$ and \hat{c} such that

$$\begin{aligned} \frac{\varepsilon}{1-\varepsilon} &= H(\hat{c}) + \frac{1}{\Delta(1)} \int_{\hat{c}}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) (H(\hat{c}) - H(c)) dc, \\ \bar{\pi} &= \Delta(0) - \hat{c} - \int_{\hat{c}}^1 \frac{\bar{\pi}}{\Delta(1)} \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) dt \end{aligned}$$

if $\varepsilon \in (\varepsilon'', \bar{\varepsilon})$,

- $\bar{\pi} = \bar{\pi}^{\max}$ and $\hat{c} = 0$ if $\varepsilon \geq \bar{\varepsilon}$,

where ε' , ε'' , and $\bar{\varepsilon}$, respectively, are defined by

$$\begin{aligned} \frac{\varepsilon'}{1-\varepsilon'} &= H(\Delta(0)), \\ \frac{\varepsilon''}{1-\varepsilon''} &= H(\Delta(0)) + \frac{1}{\Delta(1)} \int_{\Delta(0)}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) (H(\Delta(0)) - H(c)) dc, \\ \frac{\bar{\varepsilon}}{1-\bar{\varepsilon}} &= H(0) + \frac{1}{\Delta(1)} \int_0^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) (H(0) - H(c)) dc. \end{aligned}$$

PROOF. As argued above, when $\bar{\pi}^{\min} < 0$, the derivative of $\Pi_F(\bar{\pi})$ at $\bar{\pi}$ is equal to 0, implying that $\pi = \bar{\pi}^{\min}$ is optimal if and only if $\varepsilon = 0$. We also noted that $\Pi_F(\bar{\pi})$ has a corner at $\bar{\pi} = 0$. The threshold ε' is the value of ε at which the slope of the indifference curve is equal to the left derivative of $\Pi_F(\bar{\pi})$ at $\bar{\pi} = 0$. Recalling that for $\bar{\pi} < 0$, we have $\Pi'_F(\bar{\pi}) = -H(\Delta(0) - \bar{\pi})$, ε' is characterized by

$$-\frac{\varepsilon'}{1-\varepsilon'} = -H(\Delta(0)).$$

For all $\varepsilon < \varepsilon'$, the optimal mechanism is then characterized by the tangency point $-\frac{\varepsilon}{1-\varepsilon} = -H(\Delta(0) - \bar{\pi})$, while \hat{c} is such that the buyer's payoff at $c = 0$ is equal to $\bar{\pi}$, i.e., $\hat{c} = \Delta(0) - \bar{\pi}$.

Next, ε'' is the value of ε such that the slope of the indifference curve is equal to the right derivative of $\Pi_F(\bar{\pi})$ at $\bar{\pi} = 0$. Noting that at $\bar{\pi} = 0$, condition (5) yields $\hat{c} = \Delta(0)$; ε'' is defined by

$$-\frac{\varepsilon''}{1 - \varepsilon''} = -H(\Delta(0)) - \frac{1}{\Delta(1)} \int_{\Delta(0)}^1 \exp\left(\int_c^1 \frac{v'(t)}{\Delta(t)}\right) (H(\Delta(0)) - H(c)) dc.$$

For all $\varepsilon \in [\varepsilon', \varepsilon'']$, we then have $\bar{\pi} = 0$ and $\hat{c} = \Delta(0)$.

Finally, for $\varepsilon > \varepsilon''$, the characterization of the optimal mechanism is analogous to the mechanism in Proposition B.1 when $\varepsilon > \underline{\varepsilon}$. \square

REFERENCES

- Antić, Nemanja (2014), “Contracting with unknown technologies.” Unpublished paper, Princeton University. [177]
- Bergemann, Dirk and Karl Schlag (2011), “Robust monopoly pricing.” *Journal of Economic Theory*, 146, 2527–2543. [177]
- Bodoh-Creed, Aaron L. (2012), “Ambiguous beliefs and mechanism design.” *Games and Economic Behavior*, 75, 518–537. [178]
- Bose, Subir and Arup Daripa (2009), “A dynamic mechanism and surplus extraction under ambiguity.” *Journal of Economic Theory*, 144, 2084–2114. [178]
- Bose, Subir and Suresh Mutuswami (2012), “Bilateral bargaining in an ambiguous environment.” Working Paper No. 12/10, University of Leicester. [178]
- Bose, Subir, Emre Ozdenoren, and Andreas Pape (2006), “Optimal auctions with ambiguity.” *Theoretical Economics*, 1, 411–438. [178]
- Bose, Subir and Ludovic Renou (2014), “Mechanism design with ambiguous communication devices.” *Econometrica*, 82, 1853–1872. [178]
- Carrasco, Vinicius, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira (2017), “Optimal selling mechanisms under moment conditions.” Working paper, Vancouver School of Economics. [177, 184]
- Carrasco, Vinicius and Humberto Moreira (2013), “Robust decision making.” Unpublished paper, Fundação Getulio Vargas. [177]
- Carroll, Gabriel (2015), “Robustness and linear contracts.” *American Economic Review*, 105, 536–563. [177, 182]
- Carroll, Gabriel and Delong Meng (2016), “Locally robust contracts for moral hazard.” *Journal of Mathematical Economics*, 62, 36–51. [177]
- Condie, Scott and Jayant V. Ganguli (2017), “Pricing effects of ambiguous private information.” Working paper, Brigham Young University, SSRN 2154383. [177]
- di Tillio, Alfredo, Nenad Kos, and Matthias Messner (2017), “The design of ambiguous mechanisms.” *The Review of Economic Studies*, 84, 237–276. [178]

- Dow, James and Sérgio Ribeiro da Costa Werlang (1992), “Uncertainty aversion, risk aversion, and the optimal choice of portfolio.” *Econometrica*, 60, 197–204. [177]
- Frankel, Alexander (2014), “Aligned delegation.” *American Economic Review*, 104, 66–83. [177, 182]
- Garrett, Daniel F. (2014), “Robustness of simple menus of contracts in cost-based procurement.” *Games and Economic Behavior*, 87, 631–641. [177, 182]
- Gilboa, Itzhak and David Schmeidler (1989), “Maxmin expected utility with a non-unique prior.” *Journal of Mathematical Economics*, 18, 141–153. [176, 178]
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), “A smooth model of decision making under ambiguity.” *Econometrica*, 73, 1849–1892. [177, 191, 200]
- Kopylov, Igor (2016), “Subjective probability, confidence, and Bayesian updating.” *Economic Theory*, 62, 635–658. [184]
- Myerson, Roger B. and Mark A. Satterthwaite (1983), “Efficient mechanisms for bilateral trading.” *Journal of Economic Theory*, 29, 265–281. [178]
- Salanié, Bernard (2005), *The Economics of Contracts*. MIT Press, Cambridge, Massachusetts. [179]
- Samuelson, William (1984), “Bargaining under asymmetric information.” *Econometrica*, 52, 995–1005. [176, 177, 179, 181, 182]
- Wolitzky, Alexander (2016), “Mechanism design with maxmin agents: Theory and an application to bilateral trade.” *Theoretical Economics*, 11, 971–1004. [178]

Co-editor Ran Spiegler handled this manuscript.

Manuscript received 30 December, 2015; final version accepted 8 February, 2017; available online 16 February, 2017.