Optimal robust bilateral trade: bid-ask spread

Jernej Čopič*

May 25, 2017

Abstract

We provide an example of an optimal robust bilateral trading mechanism where there is a bid-ask spread in some trading events. The mechanism can neither be represented as, nor dominated by, a randomized posted price, contrary to the main assertions in [6] and [3]. Bid-ask spread may be (constrained) Pareto optimal in the absence of any external forces or frictions beyond incentive and participation constraints. This provides a link between posted prices, which reveal as little about demand as possible, and Vickrey-Clarke-Groves mechanisms, which are demand revealing but fail either individual rationality or budget balance. A general probability-prices representation of incentive compatible mechanisms is given by probability of trade and personalized prices. Weakened theorems in [6] and [3] are reconciled by assuming one price.

*Jernej Čopič, is at Aix-Marseille Universite, and a visiting researcher at CERGE-Ei. I am grateful to Yuming Huang for valuable assistance. I am also grateful to the seminar audience at FGV-Rio de Janeiro. Email correspondence: jernej.copic@univ-amu.fr

1 Introduction

In their ground-breaking work, [6] addressed the bargaining problem between a risk-neutral seller and buyer with private information over an indivisible item in a robust setting. Equilibrium outcomes of such bargaining are characterized by pricing mechanisms which are feasible and such that each trader has a dominant strategy. [6] argued that (randomized) posted-price mechanisms are "essentially the only (feasible) mechanisms such that each trader has a dominant strategy". Under some mild technical conditions, they then provided an apparent proof of the claim. In recent work, [3] considered the same problem of pricing mechanisms which satisfy the above desiderata and are additionally as efficient as possible – optimal.[1] They apparently showed that optimal robust pricing mechanisms are equivalent to non-wasteful randomized posted prices.

These two results seem plausible, however, as stated, both theorems are in fact untrue. Here we show that there exist optimal robust trading mechanisms where the price paid by the buyer sometimes exceeds the receipt of the seller. That is, we show existence of mechanisms which exhibit a bid-ask spread, or money burning, and can neither be represented nor dominated by randomized posted prices (or slightly more general randomized bid-ask posted prices).[2] The intuition is that in the event

[1]A trading mechanism is robust if it satisfies *ex post* incentive compatibility – which is in the present setting equivalent to dominant-strategy incentive compatibility, see [8] and [1] – balanced budget and individual rationality, and it is optimal if it satisfies an *ex post* Pareto criterion, see [2].

[2]We adopt to the term "bid-ask" spread following the terminology in [9] who uses that term for a mechanism that has similar features.

of burning some of the surplus at the expense of one of the traders, the incentive constraints are relaxed for that trader. That leads to a higher probability of trade and, consequently, a higher expected payoff to the other trader. At the minimum, then, the claims of [6] and [3] requisite some qualification: both implicitly (and perhaps unknowingly) assume the law one price.

For the characterization and optimality claims in [6] and [3] the simple assumption of the law one price seems appealing and innocuous, which, as it turns out, it is not. A deeper understanding of why the assumption matters sheds light on broader questions of economic relevance. On the technical side, the literature in mechanism design has long been accustomed to parametrizing allocations, say in an exchange, by the probability of transferring the good, and the expected transfers between the individuals (see e.g., [12] and others). In part, such reduced form derives from concerns regarding classical (or *ex post*) efficiency of allocation rules: any allocation rule where the law of one price does not hold everywhere is clearly not *ex post* efficient. Nevertheless, [12] show that in a bilateral exchange, under weaker *interim* constraints, efficient allocations are not possible due to incentive and participation constraints so that in the present robust setting such efficient allocations are not possible *a fortiori*. When classically efficient rules are not possible it is no longer clear that the most efficient allocation rules – "most efficient" defined in some appropriate sense – should satisfy the law of one price and such parametrization is no longer justified by efficiency considerations. [3]

We show that the canonical representation of incentive compatible mechanisms is by the probability of transferring the good and *personalized* expected prices (or transfers) which need not be equal. Any incentive compatible mechanism can be represented in this way, and the appropriate such representation preserves incentive compatibility. In contrast, if one were to represent an incentive compatible mechanism by the probability of transferring the good and one price, e.g., the expected price paid by one trader to the other, such representation does not necessarily preserve incentive compatibility. Indeed, in Section 3 we show that a sufficient condition for such reduced-form representation to preserve incentive compatibility is precisely that the budget balance condition holds with equality everywhere, i.e., the law of one price holds in any event in the original mechanism. In other words, it is not without loss of generality to assume the law of one price. What is perhaps more surprising is that the law of one price is not even necessary for optimality.

The flip side of this coin, when the law of one price does not hold, illuminates a strong connection to a well-known class of optimal allocation rules. Again, the example of bilateral trade serves as a useful illustration. Consider the original claim by [6] that every robust trading mechanism can be represented as a randomization over

[3][12] characterize the *ex ante* incentive efficient (optimal) allocation rules. By using the appropriate envelope theorem, see [10], one can then characterize these probability and price functions. More generally, it makes sense to consider allocation rules that are as efficient as possible under the given constraints – here, following [2] we consider the robust optimal decision rules, that is, the decision rules which cannot be Pareto dominated under any prior that the traders may hold.

posted prices. In a posted price trade is executed precisely when the seller's cost is below the price and the buyer's value is above the price. Therefore, a randomization over posted prices reveals as little information about the demand (and supply) as possible and the little information that is revealed comes from the individual rationality and budget balance constraints. In contrast, a Vickrey-Clarke-Groves scheme (VCG) is a fully demand-revealing mechanism, however, individual rationality and budget balance constraints are not simultaneously satisfied and there are indeed personalized transfers.[4] The mechanism presented here provides a link between these two extreme cases: because personalized prices are not equal, it reveals some information about the demand than the posted price, while still simultaneously satisfying all the constraints of a trading mechanism.

Besides the main purpose – to highlight the necessity of assuming one price in the aforementioned claims – the present example also suggests an explanation for a bid-ask spread that rests solely on incentive-efficiency. Most existing explanations rest on arguments of competition between market makers who intermediate trade (e.g., [5] and [4]). In contrast, the example here shows that in bilateral exchange, i.e. in a very thin market which is far from competitive, a bid-ask spread may in some instances be Pareto efficient, given the incentive and participation constraints. That is, in a thin market a bid-ask spread may arise due to informational constraints alone, in the absence of any market maker or any other intermediation whatsoever.

[4] A demand-revealing VCG scheme allocates the good efficiently but it is not classically efficient precisely because some of the numeraire must be burned in order to satisfy the incentive constraints.

In the next section we briefly describe the setting following [3], state the main results of [6] and [3], and provide our example showing that these results are not true without some additional assumptions. In the last section we show that results in [6] and [3] are true under the assumption that the price paid by the buyer equals the payment received by the seller.

2 Bid-ask spread

A seller, 1, and a buyer, 2, bargain over the price of an indivisible good, where v_1 is the seller's cost of producing the good, and v_2 is the buyer's value of the good. It is common knowledge that $v \in (0,1)^2$ and that their distribution is absolutely continuous with respect to the Lebesgue measure, but their distribution is not common knowledge;[5] v_i is private information to trader i.

A deterministic allocation is given by $(\ell, p_1, p_2) \in \{0,1\} \times (0,1)^2$, $\ell = 1$ if there is trade buyer, $\ell = 0$ if there is no trade and p_i is the payment received or made by trader i – in general these two payments may be different. Traders have quasilinear utility functions: for the seller, $u_1(\ell, p_1, v_1) = p_1 - \ell v_1$, and for the buyer, $u_2(\ell, p_2, v_2) = \ell v_2 - p_2$.

A direct revelation mechanism (a mechanism) is a mapping μ from traders' re-

[5] Traders may have different beliefs about the type distribution, and different beliefs about each other's beliefs of the other trader and so on, i.e., any type space is allowed, see [1].

ports $\tilde{v}_i \in (0,1)$ into randomized allocations,

$$\mu : (0,1)^2 \to \Delta\left(\{0,1\} \times (0,1)^2\right)$$

That is, $\mu[\tilde{v}]$ is a lottery faced by the traders when they report $(\tilde{v}_1, \tilde{v}_2)$; denote by $E_{\mu[\tilde{v}]}$ the expectation operator with respect to the probability measure $\mu[\tilde{v}]$. Throughout, j denotes trader other than i.

Denote by $U_i^\mu(v; \tilde{v}_i)$ the (expected) payoff to trader i in a mechanism μ when traders' valuations are $v \in (0,1)^2$, trader i reports $\tilde{v}_i$, and trader j reports v_j,

$$U_i^\mu(v; \tilde{v}_i) = E_{\mu[\tilde{v}_i, v_j]} u_i(\ell, p_i, v_i)$$

Denote by $U_i^\mu(v)$ the payoff to trader i when valuations are $v \in (0,1)^2$ and both traders report their valuations truthfully.

$$U_i^\mu(v) = E_{\mu[v]} u_i(\ell, p_i, v_i).$$

A mechanism μ is *ex post* incentive compatible, if reporting valuations truthfully is an *ex post* Nash equilibrium.[6]

$$U_i^{\mu}(v) \geq U_i^{\mu}(v; \tilde{v}_i), \forall \tilde{v}_i, \forall v_i, \forall v_j, i \in \{1, 2\} \quad (1)$$

A mechanism μ satisfies *ex post* budget balance and individual rationality, if,

$$support(\mu[v]) \subset \{(\ell, p_1, p_2) \mid v_1 \times \ell \leq p_1 \leq p_2 \leq v_2 \times \ell\}, \forall v \in (0,1)^2 .$$

By *ex post* individual rationality and budget balance, we can let the state space of the realizations of the lottery $\mu[v]$ be $\Omega = [0,1)^2$, for every profile of reports v. The state space Ω then represents all possible prices, and any prices, s.t., $p_1 = 0$ signify no trade; observe that if $p_2 = 0$, then $p_1 = 0$ by budget balance.

Definition 1. A mechanism μ is a robust trading mechanism if it satisfies *ex post* budget balance, individual rationality, and incentive compatibility.

Example 1. Randomized posted price. Let $\hat{p}$ be a realization of some given random variable with the range $(0,1)$. If the traders' reports are such that $\tilde{v}_1 \leq \hat{p} \leq \tilde{v}_2$ then the two traders trade at price $\hat{p}$, and otherwise there is no trade. Note that the receipt of the seller here equals the price paid by the buyer. A different interpretation of

[6]In a separable environment, *ex post* incentive compatibility is equivalent to requiring that the trading mechanism is *interim* (or Bayesian) incentive compatible on any type space, and in particular, for any common prior distribution over payoff types F, see [1] and also [8]. The revelation principle holds (see, e.g., [11]) so that the restriction to direct revelation mechanisms is without loss of generality. Note that in a private-values environment considered here, *ex post* incentive compatibility is equivalent to dominant-strategy incentive compatibility, i.e., strategy proofness.

this mechanism is that the price is determined by some exogenous (possibly random) rule and the two traders trade if they both agree to the price. It is evident that such a pricing mechanism is robust.

Example 2. A randomized posted bid-ask price (or simply bid-ask price). Let $(\hat{p}_1, \hat{p}_2)$ be a realization of a given random variable with the range $\{(0,1)^2 \mid p_2 > p_1\}$. Here traders trade iff $\tilde{v}_1 \leq \hat{p}_1 < \hat{p}_2 \leq \tilde{v}_2$, with the interpretation that the two prices which satisfy budget balance are determined by the exogenous random rule. Note that such a pricing mechanism is robust but seems evidently "less efficient" than some appropriately defined randomized posted price, e.g., $\hat{p} \equiv \hat{p}_2$.

Two mechanisms μ, μ' are *payoff equivalent*, if,

$$U_i^{\mu}(v) = U_i^{\mu'}(v), \forall v \in (0,1)^2 .$$

A mechanism μ is a family of lotteries over allocations – in examples 1 and 2 where these are conditional probability distributions arising from some underlying probability distribution – which are intractable objects to work with. It is common practice to instead consider a simpler payoff-equivalent class of reduced-form mechanisms, which are specified by the probability of trade and the expected price, or the expected price conditional on trade taking place. That is, one can formally define a representation of a mechanism μ by the probability of trade $\hat{\varphi}$ and the price conditional on trade $\hat{\pi}$. These two functions are given as solutions to the two equations,

$$U_i^\mu(v) = \hat{\varphi}(v) u_i \left(1, \hat{\pi}_i(v), v_i\right), i = 1, 2. \tag{2}$$

When μ is robust, one can easily show that $U_i^\mu(v)$ is weakly decreasing in v_1 and weakly increasing in v_2 so that these two functions are well defined and the corresponding rule given by $(\hat{\varphi}, \hat{\pi})$ is payoff equivalent to μ.

The main results of [6], Corollaries 1-3 to Theorem 1, can now be formulated as follows.

Let the mechanism $(\hat{\varphi}, \hat{\pi})$ be robust.

1. If $(\hat{\varphi}, \hat{\pi})$ are twice-differentiable, then μ is payoff equivalent to a randomized posted price.
2. If $\hat{\varphi}$ maps onto $\{0, 1\}$, then μ is payoff equivalent to a randomized posted price.
3. If $\hat{\varphi}$ takes finitely many different values on a finite grid, then μ is payoff equivalent to a randomized posted price.

[6] then conjecture that the results are true in general and conclude that any robust trading mechanism is payoff equivalent to a randomized posted price.

Example 3. Consider the robust trading mechanism specified in Figure 1, given by the probability function φ and the two price functions π_1 and π_2. In the upper left region, $\varphi(v) = 1$ and $\pi_1(v) = \pi_2(v) = p^*$; In the lower left region, $\varphi(v) = \underline{\varphi}$ and $\pi_1(v) = \pi_2(v) = \underline{p}$; In the upper right region $\varphi(v) = \bar{\varphi}$, $\pi_1(v) = \tilde{p}$ and $\pi_2(v) = \bar{p}$, where $\bar{p} > \tilde{p}$. Further suppose that these quantities satisfy,

$$p^* = \bar{\varphi}\tilde{p} + (1 - \bar{\varphi})\bar{p} \tag{3}$$

$$p^* = \underline{\varphi}\underline{p} + (1 - \underline{\varphi})\bar{p} \tag{4}$$

It is evident that (φ, π_1, π_2) satisfies individual rationality and budget balance. We now show that it is also incentive compatible. First, it is clear that a seller of any type v_1 would not want to report a $\tilde{v}_1 < v_1$. Next, a type v_1 in the lower left region or in the upper right region would not have any incentive to report a $\tilde{v}_1 > v_1$: in the upper right region, when $\tilde{v}_1 \leq \tilde{p}$ such mis-reporting would make no difference, and when $\tilde{v}_1 > \tilde{p}$ such mis-reporting would result in a payoff 0 rather than a strictly positive payoff $(v_1 - \tilde{p})\bar{\varphi}$. Types in the upper left region would also not have any incentive to mis-report as the marginal type $v_1 = \underline{p}$ is indifferent between reporting truthfully and misreporting to some $\tilde{v}_1 \in (\underline{p}, \tilde{p}]$, that is, $p^* - \underline{p} = \bar{\varphi}(v)(\tilde{p} - \underline{p})$, which follows from (3); all other types in that region would strictly prefer to report truthfully. Incentive compatibility for the buyer follows by a similar argument from (4).

Since $\tilde{p} < \bar{p}$, (3) and (4) imply that $\underline{\varphi} + \bar{\varphi} > 1$. This implies that (φ, π_1, π_2) is not payoff equivalent either to a randomized posted price, or to a randomized bid-ask price. This is at odds with the main claim in [6].

It may be surmised that the mechanism from Example 3 rests on some inherent inefficiency due to the bid-ask spread so that it bears no practical relevance. More precisely, is the mechanism in Example 3 less efficient (in the Pareto sense) than some randomized posted price?

[3] address the efficiency of robust trading mechanisms. To do so they first prove that the three corollaries in [6] can indeed be restated as a more general result and it is not necessary to impose any technical conditions on $(\hat{\varphi}, \hat{\pi})$ for the characterization to hold. [3] then apply a suitable notion of robust Pareto efficiency, defined in [2].

Given two robust trading mechanisms μ and μ', μ' (*ex post*) Pareto dominates μ if, $U_i^{\mu'}(v) \geq U_i^{\mu}(v)$, $\forall v \in (0,1)$, $i \in \{1,2\}$, and there is a trader $i \in \{1,2\}$ and a subset of types $O \subset (0,1)^2$, such that the Lebesgue measure of O is positive and,

$$U_i^{\mu'}(v) > U_i^{\mu}(v), \forall v \in O.$$

Therefore, a mechanism μ dominates μ' if μ yields a strictly higher utility to at least one of the traders in some event (that has a positive probability under any joint distribution over types with a full support); and μ is no worse for either trader in any event.[7]

Definition 2. A robust trading mechanism μ is *optimal* if there does not exist a robust trading mechanism μ' that Pareto dominates μ.

In a randomized posted price some non-zero probability might be assigned to prices at which at least one of the traders would in no event be willing to trade. Čopič and Ponsatí define a randomized posted price to be *non wasteful* if that is not the case. Their main result, Theorem 1, is as follows,

Theorem (Čopič and Ponsatí 2016) A mechanism is an optimal robust trading mechanism if and only if it is payoff equivalent to a non-wasteful randomized posted price.

Consider again the mechanism from Example 3. It seems plausible that (φ, π_1, π_2) should be dominated by some randomized posted price. However, that is not the case: (φ, π_1, π_2) is not dominated by any randomized posted price. Suppose to the contrary. In the upper left region, the price has to coincide with p^* and the probability of trade must be 1. By the incentive constraints of the buyer the price and the probability of trade also cannot change in the lower left region. Therefore, the only possible

[7]In the definition of optimality it is key that the mechanism dominating a given robust mechanism is also robust. That is, the dominating mechanism must be feasible according to the desiderata dictated by a given allocation problem. For instance, a randomized bid-ask price defined in Example 2 is not an optimal robust mechanism – it is dominated by some appropriately defined randomized posted price. In a Bayesian setting, [7] define a closely related notion of *ex post* incentive efficiency and then argue that in such settings classical Pareto efficiency should be used instead, with no regards for informational constraints. However, classically Pareto efficient mechanisms may fail to exist and the above notion of constrained Pareto optimality is an appropriate notion for robust mechanism design. See [2] for more details.

candidate for a dominating randomized posted price is to draw the price $\underline{p}$ with probability $1-\bar{\varphi}$ and the price $\bar{p}$ with probability $\bar{\varphi}$. It can easily be shown that under such a randomized posted price, all types of the seller in the upper right region are better off. Additionally some types, which do not trade under (φ, π_1, π_2), trade under this randomized price. Therefore, all seller's types are better off under such a randomized posted price. However, since $1-\bar{\varphi} < \underline{\varphi}$, under the randomized posted price, buyer's types in the upper right region trade with a lower probability (and at the same price as before). These types of the buyer are strictly worse off than under (φ, π_1, π_2). Thus, (φ, π_1, π_2) is not dominated by any randomized posted price, which is at odds with the main result in [3].

Example 3 thus shows that not only do there exist mechanisms with a bid-ask spread, which are not payoff equivalent to randomized posted prices (or randomized bid-ask prices), but that such mechanisms may even be optimal. This mechanism has a simple and familiar intuition: burning some of the surplus at the expense of some types of one of the traders relaxes the incentive constraints and thus favors some types of the other trader by increasing the probability of trade. By way of the bid-ask spread which depends on traders' private information, some of the sellers' types are excluded from trading while the probability of trade is in some events increased and the buyer receives a higher utility in those events. A bid-ask spread diminishes the set of events where the probability of trade is positive but increases the probability of trade on that set. This trading mechanism cannot be characterized by a randomized posted price mechanism, or a randomization over "naïve" bid-ask price mechanisms and is not Pareto dominated by a randomized posted price – it is not characterized by a lottery over mechanisms of the form (p_1, p_2), where $p_1 \leq p_2$ and the traders trade if the seller agrees to the receipt of p_1 and the buyer agrees to the payment of p_2. This implies existence of optimal robust trading mechanisms other than randomized posted prices. What goes wrong, or alternatively, what additional assumption must be made in to turn the assertions in [6] and [3] into actual theorems?

3 Discussion

The key issue with the above assertions is payoff equivalence: if two mechanisms are payoff equivalent, incentive compatibility of one does not imply incentive compatibility of the other. Incentive compatibility of a mechanism implies incentive compatibility of a payoff-equivalent mechanism under the assumption of the law of one price. With that additional assumption, as a corollary, the claims by [6] and [3] hold.[8] More generally, incentive compatibility carries over between two payoff-equivalent mechanisms if the two have the same canonical *probability-prices represenation.*

Given a robust trading mechanism μ, let $\varphi(v) = E_{\mu[v]}\ell = \mu[v](\{\ell = 1\})$, and $\tilde{\pi}_i(v) = E_{\mu[v]}p_i$, so that $\varphi(v)$ is the probability that the good is allocated to the buyer, and $\tilde{\pi}_i(v)$ is the expected price faced by trader i. Note that by *ex post* budget

[8] Still, both claims are significantly less powerful, e.g., the correct statement of the [3] result is that *among* all robust trading mechanisms satisfying the law of one price, all non-wasteful randomized posted prices are the optimal ones.

balance and individual rationality, the prices can only be positive whenever the object is allocated so that $\tilde{\pi}_i(v) = \varphi(v)\pi_i(v)$, where $\pi_i(v)$ is the expected price conditional on trade taking place,

$$\pi_i(v) = \begin{cases} E_{\mu[v]}(p_i \mid \ell = 1), & \text{if } \mu(\{\ell = 1\}) > 0, \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

We can therefore write,

$$U_i^{\mu}(v) = \mu[v](\{\ell = 1\})u_i\left(1, \pi_i(v), v_i\right) = \varphi(v)u_i\left(1, \pi_i(v), v_i\right),$$

or,

$$U_1^{\mu}(v) = \varphi(v)\left(\pi_1(v) - v_1\right), \text{ and, } U_2^{\mu}(v) = \varphi(v)\left(v_2 - \pi_2(v)\right), \forall v \in (0,1)^2.$$

The mechanism μ is payoff equivalent to (φ, π_1, π_2). Call (φ, π_1, π_2) the *probability-prices representation* of μ. Here, by budget balance, $\pi_1(v) \leq \pi_2(v), \forall v$. More generally, when individual rationality and budget balance are not imposed, the probability-prices representation of a mechanism is given by $(\varphi, \tilde{\pi}_1, \tilde{\pi}_2)$. In the special case, when $\tilde{\pi}_1 \equiv \tilde{\pi}_2$ denote that one price by $\tilde{\pi}$ (or π). We call such a probability-prices representation *unitary*. It is evident that under individual rationality and budget balance, in order for a mechanism to have a unitary probability-prices representation it must be that $p_1 = p_2$, almost everywhere on $\mu[v]$, for every v. A different way to state this is that the budget constraint is satisfied with equality.

The probability-prices representation of a mechanism is the canonical representation: it preserves all the properties of the original mechanism.

Theorem 1. A mechanism μ is incentive compatible, if and only if, the probability-prices representation of μ is incentive compatible.

Proof. The proof follows directly from the definition of incentive compatibility and the probability-prices representation. $\square$

Proposition 1. Let mechanisms μ and μ' have a unitary probability-prices representation. Then μ and μ' are payoff equivalent if and only if μ and μ' have the same probability-prices representation.

Proof. If μ and μ' have the same probability-prices representation then they are evidently payoff equivalent. For the converse, first note that by payoff equivalence, $U_i^{\mu}(v) \neq 0 \iff U_i^{\mu'}(v) \neq 0$ and $\varphi > 0 \iff \varphi' > 0$. By payoff equivalence, and summing over both traders, we have,

$$\varphi(v)(v_2 - \tilde{\pi}_2(v) + \tilde{\pi}_1(v) - v_1) = \varphi'(v)(v_2 - \tilde{\pi}'_2(v) + \tilde{\pi}'_1(v) - v_1),$$

and since $\tilde{\pi}_2(v) = \tilde{\pi}_1(v) = \tilde{\pi}(v)$ and $\tilde{\pi}'_2(v) = \tilde{\pi}'_1(v) = \tilde{\pi}'(v)$, it follows that $\varphi \equiv \varphi'$. Applying payoff equivalence for each trader, it follows that $\tilde{\pi} \equiv \tilde{\pi}'$. $\square$

The key for the claims in [6] and [3] is that when the probability-prices representation of the mechanism μ is unitary, incentive compatibility carries over under payoff-equivalence. This is established in the following corollary to the two above results.

Proposition 2. Let μ be a robust trading mechanism with a unitary probability-prices representation, let μ' satisfy individual rationality and budget balance and such that μ' is payoff equivalent to μ. Then μ' is a robust trading mechanism.

Proof. Individual rationality and budget balance of μ' follow from payoff equivalence. Incentive compatibility follows from propositions 1 and 1. $\square$

Even when $\pi_1 \not\equiv \pi_2$ one can define a representation of the mechanism μ with a single price $\hat{\pi}$ and the probability function $\hat{\varphi}$ as in (2). The problem is that such a representation ignores the information given by the structure of the state space of events where traders trade and at what prices. In contrast, this information is not ignored in the probability-prices representation (5). More specifically, the problem is that the original mechanism μ might be incentive compatible, but the payoff equivalent mechanism $(\hat{\varphi}, \hat{\pi})$ might not. That is precisely the case with the mechanism in Example 3: it is incentive compatible, however, the payoff equivalent mechanism of the form $(\hat{\varphi}, \hat{\pi})$ is not incentive compatible. This is why that mechanism is missed by the characterizations in [6] and [3] who assume incentive compatibility of $(\hat{\varphi}, \hat{\pi})$ rather than of the primitive mechanism μ.

In Proposition 2 it is crucial that the probability-prices representation of μ is unitary, which is the key assumption in Proposition 1. Without that assumption, mechanism μ may be incentive compatible, while a payoff-equivalent mechanism μ' might not satisfy incentive compatibility. In particular, if one formally restricts attention to robust trading mechanisms of the form $\hat{\pi}$, $\hat{\varphi}$, then one is effectively restricting attention to robust mechanisms that have unitary probability-prices representation, or, to mechanisms where the payment by the buyer equals the receipt by the seller. In [6] this assumption is made implicitly in the representation, while in [3] it is implicit in their Lemma 2. Both [6] and [3] thus assume that the pricing mechanisms have a unitary probability-prices representation. As shown in Example 3, without that additional assumption, there exist optimal robust mechanisms which are not payoff equivalent to randomized posted prices.

As a final note, we emphasize that the mechanism in the present example is quite different from the double auction with many buyers and sellers proposed by [9], where a bid-ask spread also arises in some events. The pricing in [9] for the pairs of traders who trade is akin to a bid-ask price (where the bid and the ask are determined by the valuations of the last pair traders on the supply and demand side who do not trade). McAfee's mechanism is indeed much more closely related to a posted price, or what we call a bid-ask price mechanism: the prices faced by the traders who trade do not depend on their own private information but depend only on the private information of traders who do not trade – from the perspective of the traders who trade, the mechanism is still a (randomized) bid-ask price. In contrast, the present trading mechanism aggregates more information than a randomized posted price.

References

[1] D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, 73(6):pp. 1771–1813, 2005.

[2] J. Copic. Robustness and efficiency. *mimeo*, 2015.

[3] J. Copic and C. Ponsatí. Optimal robust bilateral trade: Risk neutrality. *Journal of Economic Theory*, 163(5):276 – 287, 2016.

[4] D. Duffie, N. Garleanu, and L. H. Pedersen. Over-the-counter markets. *Econometrica*, 73(6):1815–1847, 2005.

[5] T. Gehrig. Competing markets. *European Economic Review*, 42(2):277–310, 1998.

[6] K. M. Hagerty and W. P. Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, 42:pp. 94–107, 1987.

[7] B. Holmström and R. B. Myerson. Efficient and durable decision rules with incomplete information. *Econometrica*, 51(6):pp. 1799–1819, 1983.

[8] J. O. Ledyard. Incentive compatibility and incomplete information. *Journal of Economic Theory*, 18:pp. 171–189, 1978.

[9] R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434 – 450, 1992.

[10] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

[11] R. B. Myerson. *Game Theory*. Harvard University Press, 1991.

[12] R. B. Myerson and M. A. Satterthwaite. Efficient mechanism for bilateral trading. *Journal of Economic Theory*, 29:pp. 265–281, 1983.

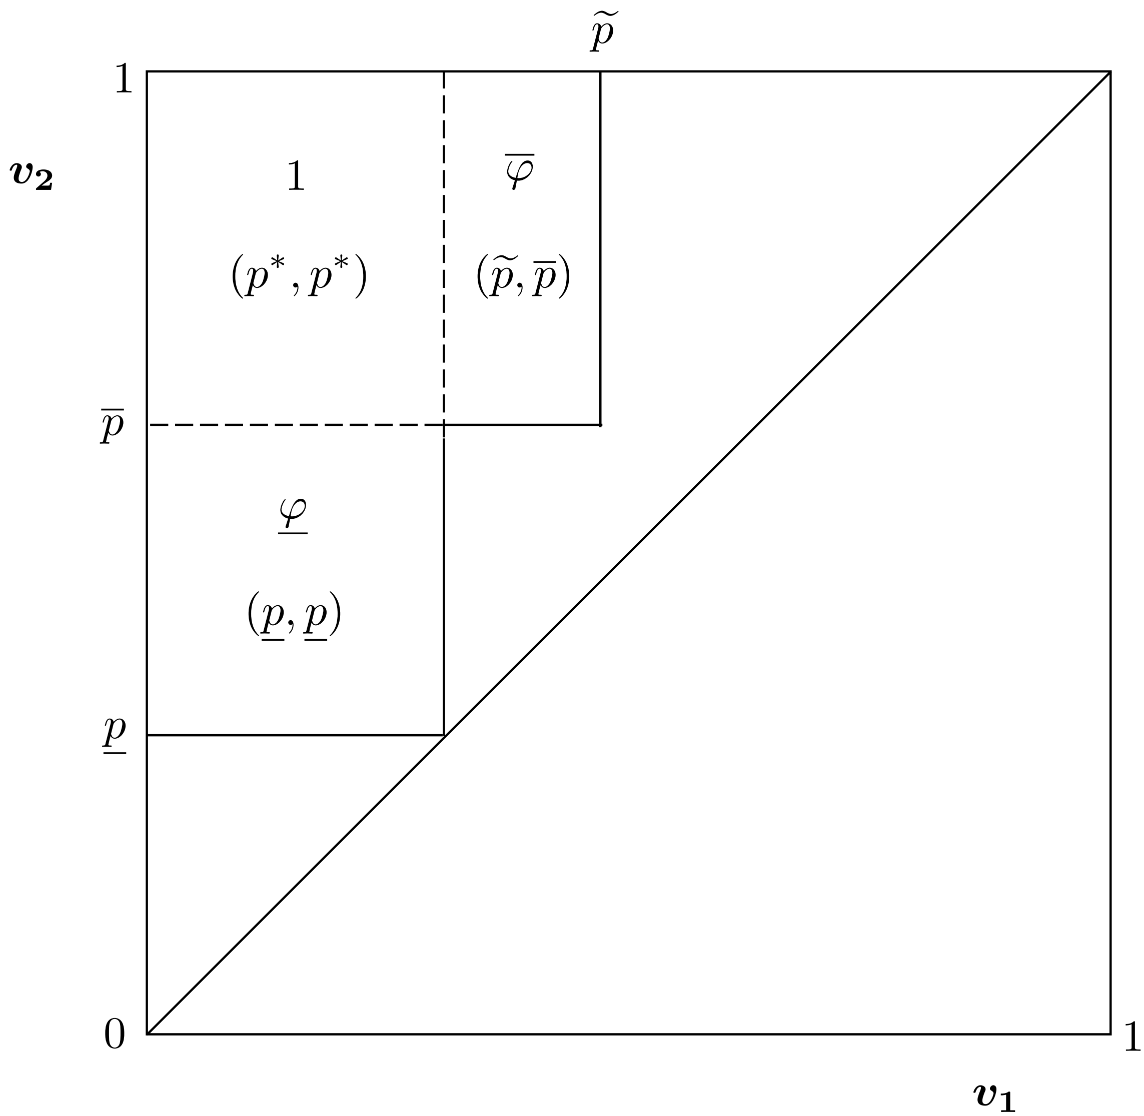
$\widetilde{p}$
1
v_2
1
(p^*, p^*)
$\overline{\varphi}$
$(\widetilde{p}, \overline{p})$
$\overline{p}$
$\underline{\varphi}$
$(\underline{p}, \underline{p})$
$\underline{p}$
0
1
v_1

Figure 1.