Waste and Efficiency in Bilateral Trade

Jernej Čopič*

August 1, 2019

Abstract

In bilateral trade, where a buyer and a seller have private cost and valuation over an indivisible good, I show that there exist constrained efficient market equilibrium outcomes that exhibit money burning: a wedge between the receipt of the seller and the price to the buyer where the remaining surplus *must be* disposed of. Previously, it has been presumed that no *ex ante* or *ex post* constrained efficient allocation exhibits such waste and that the traditional parametrization due to Myerson and Satterthwaite (1983) by probability of trade and expected payment from the buyer to the seller is sufficient to capture all market equilibrium outcomes. Such a wedge between prices cannot be represented by a lower probability of trade and the general (canonical) representation of trading outcomes is by probability of trade and personalized prices. Consequently, characterizations of market equilibrium mechanisms by posted prices as exhaustive, Hagerty and Rogerson (1987), and as constrained efficient, Copic and Ponsatí Obiols (2016), are not true without the additional assumption that there is no such wedge. Canonical representation has implications for mechanism design, bargaining, and markets.

*Jernej Čopič is at Aix-Marseille Université. I am grateful to Yuming Huang for valuable research assistance. I am also grateful to Dan Garret, Takashi Kunimoto, Humberto Moreira, Allessandro Pavan, Eric Snowberg, and the seminar audience at FGV-Rio de Janeiro. Email correspondence: jernej.copic@univ-amu.fr

"That's when I proffered my words of wisdom, that waste is the highest virtue one can achieve in advanced capitalist society,"

Haruki Murakami, Dance, dance, dance.

1 Introduction

It is fair to say that Pareto efficiency is hailed as the highest virtue by most modern economists. It is then implausible that Murakami's irreverent statement should be unconditionally true:[1] it is well known that the First Welfare Theorem holds in a general equilibrium of an exchange economy under no externalities in which case there is no waste. Nevertheless, with just one minor revision the above quote is an accurate description of the purpose of this paper:

"...that waste is *a* highest virtue one can achieve in advanced capitalist society."

I demonstrate that this is unconditionally true in a setting of a bilateral market, under circumstances where such waste has neither been thought of as plausible, and much less, virtuous. To be clear, I do not mean to question whether or not there are other highest virtues besides classical efficiency, or indeed, *constrained* efficiency –it is well known that classically efficient allocations may not always exist, see e.g., Myerson and Satterthwaite (1983). Rather, what I claim is that in a bilateral market, where it has been presumed that all constrained efficient allocations exhibit no money burning, there also exist constrained efficient allocations where some money *must* be burned, or wasted.

One of the main roles of a free market is to aggregate decentralized pieces of information held by the traders. However, when the market is thin, Myerson and Satterthwaite (1983) demonstrated

[1]Murakami might inadvertently be trespassing on rational theory – in otherwise unrelated seminal work on regulation Stiegler (1971) noted that the practice of rational theory should be the exclusive domain of economists.

that the incentive problems due to private information may also lead to impossibility of Pareto efficient market equilibrium allocations even when there are no consumption *or* information externalities. In a Bayesian setting with two risk neutral traders with quasi-linear preferences, who hold private reservation values over a single indivisible good and a common prior over these reservation values, Myerson and Satterthwaite (1983) showed that when the buyer's valuation exceeds the seller's cost – classical efficiency would then require exchange – the transaction cannot take place with certainty due to informational constraints. Myerson and Satterthwaite (1983) implicitly assumed that it was enough to consider market equilibrium outcomes where there is *no slack* in the budget balance constraint, and parametrized the equilibrium outcomes by probability of trade and expected payment from the buyer to the seller. This implicit assumption has been adopted in the literature ever since. In particular, in a scarce information setting, where traders do not have a common prior, Hagerty and Rogerson (1987) claimed that all market equilibrium mechanisms are characterized by probability distributions over posted prices; and Copic and Ponsatí Obiols (2016) claimed that all *constrained efficient* market equilibrium mechanisms are characterized by non-wasteful distributions over posted prices. Here I show that the main theorems in Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016) are incorrect, which is due to the fact that the parametrization in Myerson and Satterthwaite (1983) is not general (*canonical*). I prove that the canonical representation is by probability of trade and personalized prices, and correct the aforementioned results. The canonical representation has implications for mechanism design, bargaining, and markets.

Apart from opening the floodgates to questions regarding conditions under which a free market will guarantee efficient allocations, the Myerson and Satterthwaite (1983) impossibility result also introduced an explicit need to search for *most efficient* allocations that were theoretically possible. This necessitated a definition of *most efficient* allocations. Holmström and Myerson (1983) provided such a definition of *ex ante* and *interim incentive efficient* allocations under Bayesian incomplete information. When classical efficiency is possible, and when it is not, a key tool has been the *revelation principle*, due to Gibbard (1973), Dasgupta et al. (1979), Holmstrom (1977), and Myerson (1979): instead of studying explicit institutional arrangements, it is enough to consider the equilibrium constraints as informational constraints directly. The revelation principle leads to an enormous simplification, and in a setting where economic agents are risk neutral with quasi-linear preferences, D'Aspremont and Gerard-Varet (1979) recognized that it is enough to consider the probabilities with which different allocations obtain and the corresponding expected payments (transfers). This lead to a simple parametrization and a differential approach which allowed Myerson and Satterthwaite (1983) to prove their impossibility result. Armed with the aforementioned toolbox, Myerson and Satterthwaite (1983) also computed what was theoretically possible in an example where the distribution of these private reservation values is uniform, and noted that in the example, the allocation coincided with an equilibrium in linear strategies in a double auction due to Chatterjee and Samuelson (1983). Myerson and Satterthwaite (1983) results reverberated in the literature, and ever since, their parametrization of market equilibrium outcomes by probability of trade and expected payment from the buyer to the seller has been a quintessential tool for the study of bilateral markets and related questions, for example, the partnership dissolution problem as in Cramton et al. (1987) and McAfee (1992).

Parallel to the study of Bayesian environments was the approach under even more stringent informational constraints where the traders have no common prior – indeed the first impossibility results were under such conditions, albeit in a non-market setting, due to Arrow (1963), Gibbard (1973) and Satterthwaite (1975). Ledyard (1978) showed that in the setting with no externalities such *scarce information* implied dominant strategy equilibrium constraints, and in recent work, Čopič (2019a) defined constrained efficiency for general environments with scarce information. In the bilateral trade setting under scarce information, Hagerty and Rogerson (1987) also applied the Myerson and Satterthwaite (1983) parametrization and claimed that market equilibrium outcomes could be fully characterized as probability distributions over posted prices – a posted price is exogenous to traders' information and the traders trade if and only if both agree to the price. Copic and Ponsatí Obiols (2016) then claimed that the set of *constrained efficient* market equilibrium outcomes under scarce information was fully described by non-wasteful probability distributions over posted prices (i.e., no probability mass is assigned to prices at which one of the traders would never be willing to trade).

A key assumption in both environments, Bayesian and scarce information, has been that it is enough to consider market equilibrium outcomes where there is no slack in the budget constraint. That is, it has been assumed that the representation of market equilibrium outcomes by a probability of trade and a single (expected) payment from the buyer to the seller is general in that it represents all market equilibrium outcomes of interest. Burning money has been implicitly considered implausible – none of the market equilibrium allocations in the literature admit a wedge in the price paid by the buyer and the receipt of the seller. Moreover, in the scarce information environment there is an explicit claim that all constrained efficient outcomes exhibit *no waste*. Therefore, search as one may, but in the literature on bargaining one shall not find virtue in waste.

In the setting of scarce information, I demonstrate that a wedge in prices, where the difference must be disposed of, *is constrained efficient*.[2] It must be emphasized that the numeraire that is disposed of *cannot* in any way be redistributed between the traders – the numeraire is burned only for some profiles of reservation values, and if it were redistributed *ex post* this would ruin the traders

[2] In a different setting of an optimal contract between a principal and an agent, where there is one-sided asymmetric information, Ambrus and Egorov (2017) show that money burning can lead to an optimal contract.

incentives for truthful revelation.[3] Under scarce information this is true unconditionally, which is a direct contradiction to presumed theorems by Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016). I show that in order for these claims to be theorems, an additional assumption must be made that there is no wedge between the bid and the ask prices, that is, that the budget balance holds with equality. These theorems are then much weaker than originally claimed.

The issues thus stem from the widely-used implicit assumption that there is no slack in the budget constraint and the consequent representation of market equilibrium outcomes by probability of trade and (expected) payment from the buyer to the seller. This representation is not sufficient to capture all market equilibrium outcome mappings (*market equilibrium mechanisms*). I show that in both, the scarce information and the Bayesian environments, the sufficient or *canonical* reduced form representation is by the probability of trade and *personalized prices*. Apart from the aforementioned theorems by Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016), the canonical representation puts a question mark under the generality of some of the results in the Bayesian literature. For instance, in Cramton et al. (1987), the representation of equilibrium outcome mappings is analogous to Myerson and Satterthwaite (1983). There too, the representation with no slack in the budget balance constraint is sufficient for their main results regarding when an efficient partnership dissolution is possible, but is not general in that it would not be sufficient to study cases when efficient partnership dissolution is not possible.

The results here also suggest that in the vast literature on dynamic bargaining under incomplete information and its many applications, e.g., to search and matching, an additional and different type of inefficiency – money burning – may play a role. In dynamic bargaining, the inefficiency has

[3] Since the prior is not known, under scarce information a benevolent central planner could not redistribute a flat amount of the numeraire *ex ante*, which would not affect the incentives, even if the planner had access to frictionless insurance markets – there would be no way for the planner to acquire the appropriate insurance over *ex post* outcomes.

been represented by delay, which is equivalent to probability of trade less than one when traders' preferences are time separable. In the literature on search and matching, such probabilistic inefficiency has been described by matching frictions. However, a careful study of how money burning bilateral market equilibrium outcomes may add to these literatures is beyond the scope of this paper and is left for future research.[4] The main purpose here is to correct the theorems by Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016), amend the widely used representation of Bayesian market equilibria due to Myerson and Satterthwaite (1983), and more than anything, to show that money burning is not only necessary for a general representation of market equilibrium outcomes, but also constrained efficient.

In Section 2, I briefly describe the bilateral-trade setting and illustrate most of the intuitions by an extended example. In Section 3, I give formal definitions and Theorem 1, which states that money burning is constrained efficient. In Section 4, I provide the canonical representation of market equilibrium mechanisms and prove Theorem 1. In Section 5, I discuss the claims by Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016) and show that these are true if and only an additional assumption is made that the price paid by the buyer equals the payment received by the seller. In Section 6, I discuss the standard Bayesian environment and show that that widely-used parametrization by Myerson and Satterthwaite (1983) is not canonical and that there too the canonical representation requires personalized prices. In Section 7, the Appendix, I provide the revelation principle for the sake of completeness.

[4]Both of these literatures are vast, see e.g., the survey by Ausubel et al. (2002) on dynamic bargaining and Menzio and Trachter (2015) and Eeckhout and Kircher (2018) for some recent work on search and matching.

2 Bilateral trade and money burning

In the simplest bilateral market, there is one seller $i = 1$, one buyer $i = 2$, and one indivisible good. Traders' reservation values of the good – e.g., the seller's cost v_1 and the buyer's valuation v_2 – are private information to the traders. The range of these reservation values $V = V_1 \times V_2$ is the same for both traders and is normalized, $v_i \in V_i = (0,1)$; vectors (or profiles) are denoted in boldface, i.e., $\mathbf{v} \in V$. Furthermore, V is common knowledge, but traders' information is scarce in the sense that the traders have no knowledge of the joint distribution of reservation values, apart from the fact that all realizations are statistically possible: since the setting here is continuous, the assumption is that it is common knowledge that every open subset of reservation values $O \subset V$ has a strictly positive probability, see also Čopič (2019a). A (deterministic) final allocation is given by $(\ell, p_1, p_2) \in \{0,1\} \times \mathbb{R}^2$, where $\ell = 1$ if the good is transferred from the seller to the buyer, and $\ell = 0$ if it is not; p_i is the payment of the numeraire received or made by trader i (a positive p_i is a receipt, and a negative p_i a payment). Traders are risk neutral and their preferences are represented by separable and quasilinear utility functions: for the seller, $u_1(\ell, p_1, v_1) = p_1 - \ell v_1$, and for the buyer, $u_2(\ell, p_2, v_2) = \ell v_2 - p_2$. At each $\mathbf{v} \in V$, a market outcome may be randomized and is given by a lottery $\mu[\mathbf{v}]$ so that a market outcome mapping is given by a direct mechanism μ,

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

where $\Delta\left(\{0,1\} \times \mathbb{R}^2\right)$ denotes all lotteries over deterministic allocations. For specific traders' reservation values $\mathbf{v}$, such a lottery can be imagined as an outcome in a market, where this outcome may also entail randomness in the allocation and prices, which may be due to institutional features as well as traders' behavior. To represent a market *equilibrium* outcome mapping, certain constraints must be satisfied by a direct mechanism μ: incentive compatibility, individual rationality, and budget balance constraints. These constraints must embody the traders' lack of knowledge of probabilistic information (scarce information) which implies that these constraints are in the present setting the strongest possible, that is, point-wise or *ex post* constraints.[5] From now on, I call an incentive compatible, individually rational, and budget balanced direct revelation mechanism a *market equilibrium mechanism*. These equilibrium constraints are illustrated by the following example. More importantly, the example shows that a market equilibrium may result in money burning, and demonstrates that this market equilibrium mechanism is not Pareto dominated by a posted price. These two facts invalidate main theorems in Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016) (see detailed discussion in Section 5 below).

2.1 Example

In the figure, the seller's reservation value is on the x-axis, and the buyer's reservation value is on the y-axis. There are two outcome mappings, represented by (φ, π) and (φ', π'), where $\varphi(\mathbf{v})$ and $\varphi'(\mathbf{v})$ denote respective probabilities that a transaction takes place, and $\pi(\mathbf{v}) = (\pi_1(\mathbf{v}), \pi_2(\mathbf{v}))$ and $\pi'(\mathbf{v}) = (\pi'_1(\mathbf{v}), \pi'_2(\mathbf{v}))$ denote the respective prices faced by the two traders conditional on transaction taking place when their reported reservation values are $\mathbf{v}$ – the first price is the receipt by the seller and the second price is the payment by the buyer. Both of these outcome mappings are simple in that conditional on trade taking place, there is to each trader no variation in the possible realizations of price levels that she may face. Later I will show that such a representation is a sufficient – *canonical*

[5] Čopič (2019b) provides a lengthier discussion of scarce information and relationship with the literature on rich type spaces, see Mertens and Zamir (1985) and Brandenburger and Dekel (1993).

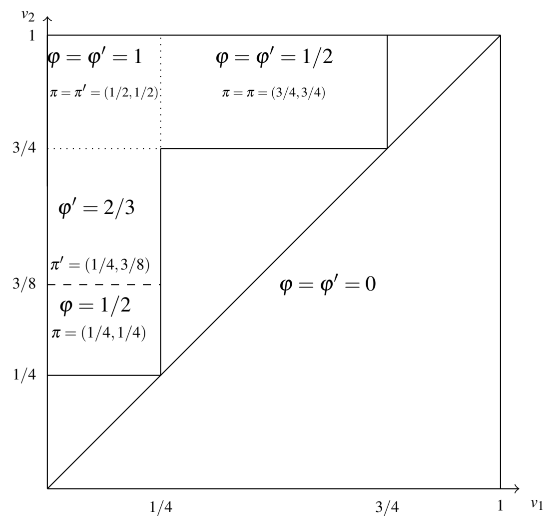

Figure 1

– reduced form to represent any possible market equilibrium mechanism μ.

I first verify that (φ, π) is a market equilibrium mechanism. Below the diagonal it must be that no transaction takes place, hence $\varphi = 0$. If that were not the case, then, either at least one of the traders would have to incur a loss (violating individual rationality or voluntary participation), or the market would have to be subsidized (violating budget balance or no subsidies). Note that $\varphi = 0$ below the diagonal is also necessary for Pareto efficiency. Now consider the area outlined by the solid line. In (φ, π), the two prices coincide everywhere in the outlined area, i.e., there is one transaction price, which implies that no subsidies are necessary in (φ, π) and budget balance is satisfied with equality. The area consists of three different regions, the upper-right region to the right of the vertical dotted line where $\varphi = \frac{1}{2}$ and the price is $\pi_1 = \pi_2 = \frac{3}{4}$, the upper-left region between the two dotted lines where $\varphi = 1$ and the price is $\frac{1}{2}$, and the lower-left region below the horizontal dotted line and between the two solid lines where $\varphi = \frac{1}{2}$ and the price is $\frac{1}{4}$. It is evident that this outcome mapping also satisfies individual rationality, for example, in the upper-right region the buyer's valuation is above and the seller's cost is below the price.

The last property that needs to be verified is incentive compatibility, that is, no trader can manipulate the allocation to her advantage by misrepresenting her reservation value. That this is sufficient for verifying the equilibrium conditions (apart from the above individual rationality and budget balance) is a consequence of the *revelation principle* due to Holmstrom (1977), Dasgupta et al. (1979), andMyerson (1979). See Section 7, the Appendix, for a formal statement. Consider the buyer (the argument for the seller is analogous). Nowhere in the parameter space is it beneficial for the buyer to misrepresent that her valuation is higher than the truth – the outcome is either unaffected or it results in a loss to the buyer. To see that it is not beneficial for the buyer to misrepresent that her valuation is lower than the truth, consider each of the three regions of the delimited area. In the upper-right and in the lower-left regions, by misrepresenting, the outcome is either unaffected or she obtains a zero utility (versus a positive expected utility if she does not misrepresent). In the upper-left region, consider the marginal buyer with the valuation $v_2 = \frac{3}{4}$. If she does not misrepresent, she obtains the good with certainty at the price $\frac{1}{2}$ so that her utility is $\frac{3}{4} - \frac{1}{2} = \frac{1}{4}$, and by pretending that her valuation is lower (but still greater than $\frac{1}{4}$) she obtains the good with a probability $\frac{1}{2}$ at the price $\frac{1}{4}$, and since her true valuation remains the same (just the outcome had changed) her utility is $\frac{1}{2}(\frac{3}{4} - \frac{1}{4}) = \frac{1}{4}$. Hence, the buyer with the marginal valuation is indifferent, which implies that for any higher valuation, her incentives for not misrepresenting are strict. Thus, regardless of what the seller's cost were – true and reported – it is optimal for the buyer to not misrepresent his information. Therefore the outcome mapping (φ, π) is an equilibrium outcome mapping. Observe that the way traders are incentivized to report truthfully is that, were a trader to misrepresent her information, the resulting more favorable price would be more than balanced out by a lower probability of trade.

Now consider the outcome mapping (φ', π'), which coincides with (φ, π) in the upper-left and the upper-right regions. The lower-left region for (φ', π') is different – it is delimited by the dashed line in the figure – and in that region, the price paid by the buyer is $\pi_2' = \frac{3}{8}$, the payment received by the seller remains the same, $\pi_1' = \frac{1}{4}$ (so that $\frac{1}{8}$ of the numeraire is disposed of), while the probability of trade in that region is $\varphi' = \frac{2}{3}$. As before, it is evident that (φ', π') satisfies individual rationality and budget balance conditions, and the incentives of the seller remain unaffected. The buyer's incentives in the upper-right region also remain unaffected, and in the lower-left region it is still not beneficial for the buyer to misrepresent her value as being lower than the truth – the marginal valuation of the buyer to still trade is $\frac{3}{8}$, which is the price at which the buyer trades in that region. For the upper-left region, consider again the buyer with the marginal valuation $\frac{3}{4}$. As before, the buyer is at that valuation indifferent between misrepresenting her valuation or not: by misrepresenting she obtains $\frac{2}{3}(\frac{3}{4} - \frac{3}{8}) = \frac{1}{4}$. Hence (φ', π') is also a market equilibrium mechanism. The key feature is that because the price in the lower-right region is now closer to that in the upper-left region, the probability of trade in the lower-left region (albeit the region is smaller than before) has risen to $\frac{2}{3}$, and that has left the buyer's incentives unaffected.

Finally, compare the payoffs to the two traders for each realization of reservation values under the two market equilibrium mechanisms (φ, π) and (φ', π'). In both, the payoffs to both traders are identical for all profiles of reservation values except in the lower-left region. Consider therefore the lower-left region. First, there is a set of traders' valuations between the dashed line at $\frac{3}{8}$ and the solid horizontal line at $\frac{1}{4}$, who trade with a probability $\frac{1}{2}$ in (φ, π) and obtain strictly positive expected utilities, but who obtain zero utility in (φ', π'). For these reservation values, both traders are worse off in (φ', π'). Next, in the remainder of the lower-left region, the buyer is also worse off in (φ', π') – at valuation $\frac{3}{4}$ she is indifferent between (φ, π) and (φ', π'), and in (φ', π') her utility decreases with a greater slope and reaches 0 at $\frac{3}{8}$. However, the seller is strictly better off in (φ', π') for any reservation values in this part of the lower-left region, as she obtains the same price but the likelihood of trade is higher. That is to say that the direct revelation mechanisms (φ, π) and (φ', π') are not (*ex post*) Pareto comparable. The main theorem of the next Section 3 is that more is true: the market equilibrium mechanism (φ', π') is as efficient as possible under the equilibrium constraints, that is, it is *constrained efficient*.

3 Constrained efficient money burning

Consider a direct mechanism $\mu : V \to \Delta(\{0,1\} \times \mathbb{R}^2)$. Denote by $\mathbb{E}_{\mu[\mathbf{v}]}$ the expectation operator with respect to the probability law $\mu[\mathbf{v}]$. Denote by $U_i^{\mu}(\mathbf{v}; v_i')$ the (expected) payoff to trader i when the profile of reservation values is $\mathbf{v} \in V$, and trader i reports v_i', while j reports her true reservation value v_j,

$$U_i^{\mu}(\mathbf{v}; v_i') = \mathbb{E}_{\mu[v_i', v_j]} u_i(\ell, p_i, v_i).$$

Denote $U_i^{\mu}(\mathbf{v}) = U_i^{\mu}(\mathbf{v}; v_i)$, that is, i's expected utility when both traders report their true information.

As explained in the example of Section 2.1, a direct mechanism represents a market equilibrium mechanism if it satisfies the following (*ex post*) constraints: (i) Incentive compatibility – (1) below – a trader should have no incentives to deviate from v_i in order to generate allocations that she would face if she reported a reservation value v_i'. (ii) Individual rationality – no trader is forced to accept a transaction that yields her a negative utility. (iii) Budget balance – a market should not require any outside subsidies, the latter two constraints are given by (2) below.

Definition 1. A direct mechanism μ is a market equilibrium mechanism if,

$$U_i^{\mu}(\mathbf{v}) \geq U_i^{\mu}(\mathbf{v}; v_i'), \forall v_i', v_i, v_j \tag{1}$$

$$support(\mu[\mathbf{v}]) \subset \{(\ell, p_1, p_2) \mid v_1 \times \ell \leq p_1 \leq p_2 \leq v_2 \times \ell\}, \forall \mathbf{v} \in V \tag{2}$$

The set of all market equilibrium mechanisms is given by $M = \{\mu \mid \mu \text{ s.t., (1) and (2)}\}$.

To compare the efficiency of various market equilibrium mechanisms, I apply the definition of *constrained efficiency* due to Čopič (2019a).

Given a $\mu, \mu' \in M$, μ' (*ex post*) Pareto dominates μ if,

$$U_i^{\mu'}(\mathbf{v}) \geq U_i^{\mu}(\mathbf{v}), \forall i, \forall \mathbf{v},$$

and the inequality is strict for at least one i on an open subset of reservation values $V' \subset V$.

In other words, a market equilibrium mechanism μ' dominates μ if, μ' makes no trader worse of for any profile of reservation values, and μ' makes at least one traders strictly better off on some non-negligible subset of profiles of reservation values.

Definition 2. A market equilibrium mechanism $\mu \in M$ is constrained efficient if there does not exist a $\mu' \in M$ that Pareto dominates μ.

The direct mechanism μ given by (φ, π_1, π_2) in Figure 2 is the general version of the money burning market equilibrium mechanism from Figure 1.

In the upper-left region, $\varphi(\mathbf{v}) = 1$ and $\pi_1(\mathbf{v}) = \pi_2(\mathbf{v}) = p^*$; In the upper-right region, $\varphi(\mathbf{v}) = \varphi'$ and $\pi_1(\mathbf{v}) = \pi_2(\mathbf{v}) = p^H$; In the lower-left region $\varphi(\mathbf{v}) = \varphi''$, $\pi_1(\mathbf{v}) = p^L$ and $\pi_2(\mathbf{v}) = p^M$, where $p^M > p^L$. Further suppose that these quantities satisfy,

$$p^* = \varphi' p^H + (1 - \varphi') p^L \tag{3}$$

$$p^* = \varphi'' p^M + (1 - \varphi'') p^H \tag{4}$$

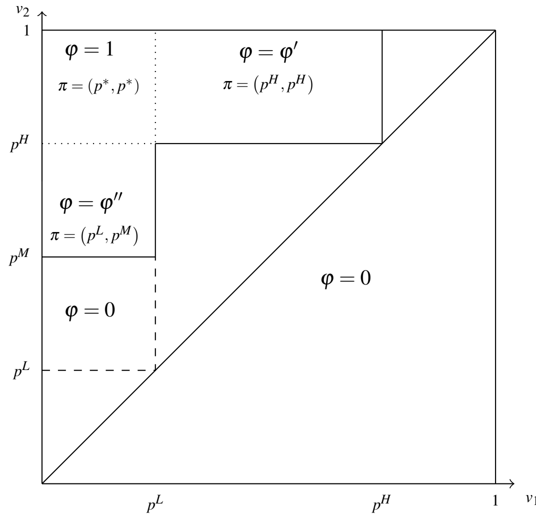

Figure 2

Since in the lower-left region, $\pi_1(\mathbf{v}) = p^L < \pi_2(\mathbf{v}) = p^M$, this direct mechanism μ exhibits money burning. That (φ, π_1, π_2) is a market equilibrium mechanism is verified as in the example of Section 2.1.

Theorem 1. The money burning market equilibrium mechanism (φ, π) depicted in Figure 2 is constrained efficient.

From example of Section 2.1 recall the intuition why the money burning market equilibrium mechanism is not dominated by a posted price: since the difference in the price for the buyer in the lower-left and upper-left regions is smaller (by way of wasting some amount of the numeraire), the probability of transaction in that region is increased for the equilibrium conditions to hold, which benefits the seller whose receipt remains unchanged at the higher probability of transaction. A detailed proof of constrained efficiency and the related discussion of canonical representation of market equilibrium mechanisms are the objectives of the next section.

4 Canonical representation of market equilibrium mechanisms and constrained efficiency of money burning

A market equilibrium mechanism implicitly describes a bilateral market institution and the outcome mapping of traders' equilibrium behavior in that institution. Since a general (possibly non-equilibrium) direct mechanism μ is given by a family of lotteries over allocations, which are generally intractable objects, it is useful to consider simpler payoff-equivalent direct mechanisms.

The traditional payoff-equivalent reduced form of a given direct mechanism is specified by the probability of trade and the expected price, or the expected price conditional on trade taking place. In the present context of bilateral trade, this representation was under the Bayesian equilibrium constraints first introduced by Myerson and Satterthwaite (1983). That is, one can formally define a representation of a direct mechanism μ by the probability $\hat{\varphi}$ that the object will be transferred from the seller to the buyer and the amount of the numeraire $\hat{\pi}$ that the buyer will then transfer to the seller. These two functions are given as the solutions to the two equations,

$$U_i^{\mu}(\mathbf{v}) = \hat{\varphi}(\mathbf{v})u_i\left(1, \hat{\pi}_i(\mathbf{v}), v_i\right), i = 1, 2. \tag{5}$$

When μ is an equilibrium outcome mapping, i.e., a market equilibrium mechanism, one can easily show that $U_i^{\mu}(\mathbf{v})$ is weakly decreasing in v_1 and weakly increasing in v_2, see e.g., Myerson and Satterthwaite (1983) for the Bayesian case and Copic and Ponsatí Obiols (2016) for the present case, so that these two functions are well defined. The corresponding direct mechanism $(\hat{\varphi}, \hat{\pi})$ is payoff equivalent to μ. A key implicit claim justifying such a reduced form is that rather than considering all market equilibrium mechanisms μ, it is enough to consider all market equilibrium mechanisms in reduced form $(\hat{\varphi}, \hat{\pi})$; an analogous observation has been made in Bayesian settings.

The problem is that this representation imposes equilibrium conditions on the reduced form directly, rather than on the general outcome mapping from which the reduced form can then be derived in such a way as to preserve the equilibrium conditions. The traditional representation thus ignores the information given by the structure of the state space of events where traders trade and at what prices. That is, the original direct mechanism μ might be incentive compatible, but the payoff equivalent mechanism $(\hat{\varphi}, \hat{\pi})$ might not. That is indeed the case with the market equilibrium mechanism of Figure 1: it is incentive compatible, however, the payoff-equivalent outcome mapping represented by the probability of trade and a single price, $(\hat{\varphi}, \hat{\pi})$, is not incentive compatible. Therefore, if one considers *only* the market equilibrium mechanisms of the form $(\hat{\varphi}, \hat{\pi})$, the market equilibrium mechanism in Figure 1 is missed.

I next formulate the *probability-prices representation of a direct mechanism* and show that it is a *canonical* representation in the sense that any direct mechanism μ may be represented and the representation preserves all the relevant equilibrium properties of the original mechanism μ.

Given an outcome mapping μ, let,

$$\varphi(\mathbf{v}) = \mathbb{E}_{\mu[\mathbf{v}]}\ell = \mu[\mathbf{v}](\{\ell = 1\}) \text{ and } \hat{\pi}_i(\mathbf{v}) = \mathbb{E}_{\mu[\mathbf{v}]}p_i, \tag{6}$$

so that $\varphi(\mathbf{v})$ is the probability that the good is allocated to the buyer, and $\hat{\pi}_i(\mathbf{v})$ is the *personalized* expected receipt (payment) of trader i. Note that when *ex post* budget balance and individual rationality are satisfied, the prices can only be positive whenever the object is allocated so that $\hat{\pi}_i(\mathbf{v}) = \varphi(\mathbf{v})\pi_i(\mathbf{v})$, where $\pi_i(\mathbf{v})$ is the expected price conditional on trade taking place,

$$\pi_i(\mathbf{v}) = \begin{cases} \mathbb{E}_{\mu[\mathbf{v}]}(p_i \mid \ell = 1), & \text{if } \mu(\{\ell = 1\}) > 0, \\ \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

We can therefore write,

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

or,

$$U_1^{\mu}(\mathbf{v}) = \varphi(\mathbf{v})(\pi_1(\mathbf{v}) - v_1), \text{ and, } U_2^{\mu}(\mathbf{v}) = \varphi(\mathbf{v})(v_2 - \pi_2(\mathbf{v})), \forall \mathbf{v}.$$

The outcome mapping μ is thus payoff equivalent to (φ, π_1, π_2). Call (φ, π_1, π_2) the *probability-prices representation* of μ. Here, by budget balance, $\pi_1(\mathbf{v}) \leq \pi_2(\mathbf{v}), \forall \mathbf{v}$. More generally, when individual rationality and budget balance are not imposed, the probability-prices representation is given by (φ, π_1, π_2). Incentive compatibility carries over between two payoff equivalent direct mechanisms if and only if the two have the same probability-prices representation.

Theorem 2. An outcome mapping μ is incentive compatible, if and only if, the probability-prices representation of μ is incentive compatible.

Proof. Suppose that μ is incentive compatible and consider the seller. Then,

$$U_1^{\mu}(\mathbf{v}; v_1') = \mathbb{E}_{\mu}[v_1', v_2] p_1 - \ell v_1 = \pi_1(v_1', v_2) - \varphi(v_1', v_2) v_1.$$

Therefore,

$$U_1^{\mu}(\mathbf{v}; v_1') \le U_1^{\mu}(\mathbf{v}) \iff \pi_1(v_1', v_2) - \varphi(v_1', v_2)v_1 \le \pi_1(\mathbf{v}) - \varphi(\mathbf{v})v_1.$$

The argument for the buyer is identical. □

Corollary 3. A direct mechanism μ represents a market equilibrium outcome, if and only if, the probability-prices representation of μ is incentive compatible, individually rational, and budget balanced.

Proof. Incentive compatibility follows from Theorem 2. The rest follows from the fact that μ satisfies individual rationality and budget balance if and only if its probability-prices representation satisfies individual rationality and budget balance. □

We are now ready to prove Theorem 1.

Proof of Theorem 1. It is evident that (φ, π_1, π_2) satisfies individual rationality and budget balance. I now show that it is also incentive compatible. Consider the seller. First, it is clear that any type v_1 would not want to report a $v_1' < v_1$. Next, types v_1 in the lower-left region or in the upper-right region do not have any incentive to report a $v_1' > v_1$: in the upper-right region, when $v_1' \le p^H$ such mis-reporting would make no difference, and when $v_1' > p^H$ such mis-reporting would result in a payoff 0 rather than a strictly positive payoff $(v_1 - p^H)\varphi'$. Types in the upper-left region also do not have any incentive to mis-report. To see this observe that the marginal type $v_1 = p^L$ is indifferent between reporting truthfully and misreporting to some $v_1' \in (p^L, p^L]$, that is, $p^* - p^L = \varphi'(\mathbf{v})(p^H - p^L)$, which follows from (3); all other types in that region strictly prefer to report truthfully. Incentive compatibility for the buyer follows by a similar argument from (4). Therefore, (φ, π_1, π_2) is an equilibrium outcome mapping, i.e., market equilibrium mechanism.

What remains to be shown is that (φ, π_1, π_2) is not Pareto dominated by any other equilibrium outcome mapping. By Theorem 2, it is enough to consider equilibrium outcome mappings in a canonical representation. Suppose to the contrary, and denote by $(\underline{\varphi}, \underline{\pi}_1, \underline{\pi}_2)$ the dominating market equilibrium mechanism. In the upper-left region, the price has to coincide with p^* and the probability of trade must be 1. Consequently, by the incentive constraints of the seller, the price and the probability of trade also cannot change in the upper-right region.

Therefore, the only possibility is to dominate (φ, π_1, π_2) in the lower-left region. Assume for a moment that $(\underline{\varphi}, \underline{\pi}_1, \underline{\pi}_2)$ is constant in that region and denote by $\underline{\varphi}', \underline{p}^L, \underline{p}^M$ the values that it takes in that region. First, it must be that $\underline{p}^L \geq p^L$ – otherwise the lower-left region would diminish and some types of traders would be strictly worse off. Second, it is not possible that $\underline{p}^L > p^L$. The reason is that if that were the case, then, by the incentive constraints of the seller, in order for the price in the upper-left region to remain p^*, it must be that $\underline{\varphi}'' < \varphi''$, so that the buyer would then be worse off in the upper-right region than under (φ, π_1, π_2). Third, it is not possible that $\underline{p}^M > p^M$. The reason is that by the incentive compatibility constraint of the buyer, $\underline{p}^M$ is the marginal buyer's valuation where trade in the lower-left region still occurs with a positive probability. Thus, if $\underline{p}^M > p^M$ then the lower-left region would again be diminished. Fourth, it is not possible that $\underline{p}^M < p^M$. In that case, by the buyer's incentive compatibility constraint (4), in order for the price to remain p^* it would have to be that $\underline{\varphi}' < \varphi'$, which implies that all types of the seller in the lower-left region would be worse off.

Finally, the proof that $(\underline{\varphi}, \underline{\pi}_1, \underline{\pi}_2)$ must be constant in the lower-left region is as follows. The seller's expected utility along the left margin of the region, the line between $(0, p^M)$ and $(0, p^H)$ must be at least $\varphi' p^L$ and the buyer's expected utility is constant and equal to $p^* p^H$ along the horizontal dotted line from $(0, p^H)$ to (p^L, p^H). These two together imply that the seller's expected utility along the line between $(0, p^M)$ and $(0, p^H)$ is constant and equal to $\varphi' p^L$. Since the seller's expected utility must be positive in the whole region, in order for the buyer's incentive constraints along the horizontal dotted line to be satisfied, and for the seller's expected utility on the left margin of the region to equal $\varphi' p^L$, it must be that the probability of trade $\underline{\varphi}$ must be constant and equal to φ' on the whole region. It is immediate to show that in order for incentive compatibility to be satisfied, a constant $\underline{\varphi}$ implies that $\underline{\pi}_1$ and $\underline{\pi}_2$ must be constant as well. □

5 Bid-ask spread and posted prices

As a consequence of applying the traditional parametrization, two previous characterizations in the literature are specious. Constrained efficiency of money burning runs contrary to the main claims in Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016). The former claimed that market equilibrium mechanisms are characterized as randomizations over posted price mechanisms (under some minor thechnical conditions) and the latter claimed that the class of constrained efficient trading mechanisms can be represented by non-wasteful randomizations over posted prices. That these claims are untrue without further qualification is a consequence of Theorem 1 and it can also be shown directly. To do so, first recall the posted prices and define naïve bid-ask prices to make clear the distinction from the money burning in Figure 2.

A randomized posted price is defined as follows. Let p be a realization of some given random variable with the range $(0,1)$. If the traders' reports are such that $v_1 \leq p \leq v_2$, then the two traders trade at price p, and otherwise there is no trade. The receipt of the seller here equals the price paid by the buyer. An explicit interpretation of this trading 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. Since each trader has a dominant strategy to agree to trade precisely when that gives her some positive surplus, this is a market equilibrium mechanism.

A randomized posted naïve bid-ask price (naïve bid-ask) is defined as follows. Let (p_1, 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 $v_1 \leq p_1 < p_2 \leq v_2$, with the interpretation that the two prices which satisfy budget balance are determined by the exogenous random rule. As with a posted price, such a pricing mechanism is a market equilibrium mechanism. However, it is evidently not constrained efficient, i.e., it is Pareto dominated by some appropriately defined randomized posted price, e.g., $p \equiv p_2$.

The central claim of Hagerty and Rogerson (1987) is as follows (Hagerty and Rogerson (1987), Corollaries 1-3 to Theorem 1).[6]

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

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."

Hagerty and Rogerson (1987) then conjecture that the results are true in general and conclude that any market equilibrium mechanism is payoff equivalent to a randomized posted price.

Naïve bid-ask trading mechanisms may seem to have been missed by this claim. While that is not entirely evident, one could assert that a naïve bid-ask price is almost identical to a posted price.

[6] Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016) refer to a market equilibrium mechanism as a robust trading mechanism.

Moreover, as pointed out, barring exogenous considerations – e.g., an intermediary with market power to extract surplus from the two traders as in Čopič et al. (2019) – a naïve bid-ask price is clearly less efficient than some adequately defined posted price so that such naive posted prices are neither plausible nor relevant.In contrast, consider again the money burning mechanism depicted in Figure 2. Since $p^L < p^M$, (3) and (4) imply that $\varphi'' + \varphi' > 1$. This in turn implies that (φ, π_1, π_2) is not payoff equivalent either to a randomized posted price, or to a randomized bid-ask price. That is clearly at odds with the main claim in Hagerty and Rogerson (1987).

Copic and Ponsatí Obiols (2016) address the efficiency of market equilibrium mechanisms. To do so they first prove that the three corollaries in Hagerty and Rogerson (1987) can indeed be restated as a more general result and it is not necessary to impose the technical conditions (1)-(3) above on $(\hat{\varphi}, \hat{\pi})$ for the characterization to hold. In a randomized posted price some non-zero probability may be assigned to prices such that at every profile of reservation values at least one of the traders is in not willing to trade. If that is not the case, Copic and Ponsatí Obiols (2016) define such a randomized posted price to be *non wasteful*. Their main claim is then formulated as follows (Theorem 1 in Copic and Ponsatí Obiols (2016)):

“A direct mechanism is a constrained efficient robust trading mechanism if and only if it is payoff equivalent to a non-wasteful randomized posted price."

As shown above, the market equilibrium mechanism in Figure 2 cannot be represented by a lottery over posted prices and by Theorem 1 it is constrained efficient. In particular, it is not dominated by a lottery over posted prices. That is at odds with the main claim in Copic and Ponsatí Obiols (2016).[7]

[7]This can also be shown directly. Suppose to the contrary. In the upper-left region, the price has to coincide with p^*

The question is, what additional assumptions must be made to reconcile these claims with Theorem 1. Given that these characterizations are based on the reduced-form representation in Myerson and Satterthwaite (1983), a slightly broader question is under what conditions the traditional representation of market equilibrium mechanisms by probability of trade and one expected price is without loss of generality. That is, under what conditions is the traditional reduced-form representation incentive compatible if and only if the original outcome mapping μ is incentive compatible.

In the traditional representation, there is no money burning, i.e., $\pi_1 \equiv \pi_2$, so denote that one price by π. I call such a probability-prices representation *unitary*. Under individual rationality and budget balance, in order for an outcome mapping μ to have a unitary probability-prices representation, it must be that $p_1 = p_2$ almost everywhere on $\mu[\mathbf{v}]$, for every $\mathbf{v}$. A different way to state this is that the budget constraint must hold with equality. In the following two propositions I show that is the necessary and sufficient condition for the traditional reduced-form representation to preserve incentive compatibility.

Proposition 1. Let mechanisms μ and μ' have a unitary probability-prices representation. Then μ and μ' are equilibrium 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, $U_i^{\mu}(\mathbf{v}) \neq 0 \iff U_i^{\mu'}(\mathbf{v}) \neq 0$ and

and the probability of trade must be 1. By the incentive constraints of the seller the price and the probability of trade also cannot change in the upper-right region. Therefore, the only possible candidate for a dominating randomized posted price is to draw the price p^L with probability $1-\varphi'$ and the price p^H with probability φ'. As in the example of Section 2.1 one can verify that under such a randomized posted price, all types of the buyer in the lower right region are better off. Additionally some types who do not trade under (φ, π_1, π_2) would then trade. Therefore, all buyer's types are weakly better off under this randomized posted price, and some are strictly better off. However, under such a randomized posted price the seller's types in the lower-left region would trade at the same price and with a lower probability $1-\varphi'$ as opposed to $\varphi'' > 1-\varphi'$ under (φ, π_1, π_2). Consequently, these types of the seller are strictly worse off than under (φ, π_1, π_2) so that (φ, π_1, π_2) is not dominated by any randomized posted price.

$\varphi > 0 \iff \varphi' > 0$. By payoff equivalence, and summing over both traders, we have,

$$\varphi(\mathbf{v})(v_2 - \pi_2(\mathbf{v}) + \pi_1(\mathbf{v}) - v_1) = \varphi'(\mathbf{v})(v_2 - \pi_2'(\mathbf{v}) + \pi_1'(\mathbf{v}) - v_1),$$

and since $\pi_2(\mathbf{v}) = \pi_1(\mathbf{v}) = \pi(\mathbf{v})$ and $\pi_2'(\mathbf{v}) = \pi_1'(\mathbf{v}) = \pi'(\mathbf{v})$, it follows that $\varphi \equiv \varphi'$. Applying payoff equivalence for each trader, it follows that $\pi \equiv \pi'$. □

Thus, the characterizations of Hagerty and Rogerson (1987) and Copic and Ponsatí Obiols (2016) are true precisely under the assumption that the probability-prices representation of the equilibrium outcome mapping is unitary, that is, that there is no money burning. Such characterizations are clearly much weaker: they characterize market equilibrium mechanisms where there is no disposal of the numeraire. It should be noted that the non-wasteful probability distributions over posted prices are constrained efficient (they cannot be dominated by a bid-ask trading mechanism). However, as shown by the canonical representation, disposal of the numeraire in a bid-ask trading mechanism in Figure 2 is a property of equilibrium trading outcomes and is also constrained efficient.

6 Bayesian environments

In a Bayesian environment there is a probability law $\mathbb{P} \in \Delta(V)$, which describes the traders' common prior belief regarding the joint distribution of reservation values – thus, in addition to the description of the bilateral trade environment with scarce information (traders' reservation values V, the space of allocations, $\{0,1\} \times \mathbb{R}^2$, and utility functions $\mathbf{u}$), there is a common prior $\mathbb{P}$. As before, a direct mechanism is given by $\mu : V \to \{0,1\} \times \mathbb{R}^2$, and for a given profile of reservation values $\mathbf{v}$ and reports

v'_i, traders' expected utilities in μ are given by $U_i^{\mu}(\mathbf{v}; v'_i)$, and $U_i^{\mu}(\mathbf{v})$ when i reports truthfully. Denote,

$$\bar{U}_i^{\mu}(v_i; v'_i) = \mathbb{E}_{\mathbb{P}|v_i}\left[U_i^{\mu}(v_i, \omega_j; v'_i)\right] \text{ and } \bar{U}_i^{\mu}(v_i) = \mathbb{E}_{\mathbb{P}|v_i}\left[U_i^{\mu}(v_i, \omega_j)\right].$$

Similarly, given a μ, define as before φ, π_1, π_2, and denote,

$$\bar{\varphi}_i(v_i) = \mathbb{E}_{\mathbb{P}|v_i}\left[\varphi(v_i, \omega_j)\right] \text{ and } \bar{\pi}_{k,i}(v_i) = \mathbb{E}_{\mathbb{P}|v_i}\left[\pi_k(v_i, \omega_j)\right], \forall i, k, \text{ and } j \neq i.$$

Finally, recall that: (i) a μ is *interim* (Bayesian) incentive compatible if,

$$\bar{U}_i^{\mu}(v_i) \geq \bar{U}_i^{\mu}(v_i; v'_i), \forall i, \forall v_i, \forall v'_i, \tag{8}$$

and (ii), *interim* individually rational if,

$$\bar{U}_i^{\mu}(v_i) \geq 0, \forall i, \forall v_i, \forall v'_i. \tag{9}$$

As in Myerson and Satterthwaite (1983), define *ex ante* balanced budget (or no subsides) by,

$$\mathbb{E}_{\mathbb{P}}\left[\mathbb{E}_{\mu[\mathbf{v}]} p_1 - p_2\right] \leq 0 \tag{10}$$

A direct mechanism, which is Bayesian incentive compatible and individually rational, and *ex ante* budget balanced direct mechanism μ is called a Bayesian market equilibrium mechanism.

As in the scarce information environment, define the probability-prices representation of μ by (φ, π_1, π_2). Recall that the probability-prices representation is unitary when $\pi_1 = \pi_2 = \pi$, a.e.

on V.Observe that the probability-prices representation is equivalent to a representation given by $(\varphi, \varphi * \pi_1, \varphi * \pi_2)$, and in the unitary case to $(\varphi, \varphi * \pi)$, which is the representation in Myerson and Satterthwaite (1983).

Theorem 4. A direct mechanism μ is *interim* incentive compatible, if and only if, the probability-prices representation of μ is *interim* incentive compatible.

Proof. Consider the seller, and observe that,

$$\bar{U}_1^{\mu}(v_1; v_1') = \mathbb{E}_{\mathbb{P}|v_1}\left[\mathbb{E}_{\mu[v_1',v_2]}[p_1 - \ell v_1]\right] = \mathbb{E}_{\mathbb{P}|v_1}\left[\varphi(v_1',v_2)(\pi_1(v_1',v_2) - v_1)\right].$$

The claim for the seller follows, and a parallel argument proves the claim for the buyer. □

Corollary 5. A direct mechanism μ is an Bayesian market equilibrium mechanism, if and only if, the the probability-prices representation of μ is *interim* incentive compatible, individually rational and budget balanced.

Proof. First observe that, e.g., for the seller, $\bar{U}_i^{\mu} \geq 0$ if and only if $\mathbb{E}_{\mathbb{P}|v_1}[\varphi(v_1',v_2)(\pi_1(v_1',v_2) - v_1)]$, so that μ satisfies *interim* individual rationality – a parallel argument works for the buyer – if and only if the probability-prices representation of μ satisfies *interim* individual rationality. Similarly for budget balance. □

To prove their impossibility theorem, Myerson and Satterthwaite (1983) represent a trading mechanism by a probability of trade and the expected *unitary* price paid by the buyer to the seller. This representation is sufficient to prove the impossibility theorem – a Bayesian market equilibrium mechanism with a *unitary* probability prices representation is evidently not *ex post* Pareto efficient as some numeraire is disposed of. However, their representation is not sufficient to characterize a

Bayesian market equilibrium mechanism which maximizes the total gains from trade – See Theorem 2, p 274 in Myerson and Satterthwaite (1983) and the ensuing example. In order to have a complete characterization, non-unitary Bayesian market equilibrium mechanisms should be considered as well.More broadly, the unitary probability prices representation of Bayesian market equilibrium mechanisms has been widely accepted in the Bayesian mechanism design literature on bilateral trade and related environments as the general sufficient (canonical) representation. The above results suggests a revision of where in the literature that might or might not matter.[8]

7 Appendix: explicit markets and the revelation principle

An explicit bilateral market is defined by $(V, O, S, \gamma, \tilde{u})$, where S is the set of players' type-contingent trading strategies (possibly mixed), so that an $s \in S$ is a mapping $s : V \to O$, $s_i(\theta) = s_i(\theta_i)$; O is the outcome space, γ is the outcome mapping, $\gamma : O \to \Delta\left(\{0,1\} \times \mathbb{R}^2\right)$, and $\tilde{u} : S \times V \to \mathbb{R}^2$ are the players' indirect utility functions, $\tilde{u}(s, v) = u(\gamma(s), v), s \in S$. At any outcome realization in a voluntary and free market institution the individual rationality and budget balance conditions must hold. Then, a strategy profile s is an *ex post* Nash equilibrium if,

$$\tilde{u}_i(s, v) \geq \tilde{u}_i(s'_i, s_{-i}, v), \forall s'_i \in S_i, \forall \mathbf{v} \in V.$$

The revelation principle then states that if there exists an *ex post* Nash equilibrium s of $(V, O, S, \gamma, \tilde{u})$, then there exists an *ex post* incentive compatible direct revelation mechanism μ, such that $\mu[\mathbf{v}] =$

[8]Budget balance, a key consideration for canonical representation, was first studied by Green and Laffont (1979). A much larger part of the recent literature has been focused on how and when efficient market equilibrium outcomes may obtain, especially the applications of revenue equivalence theorem, see for instance Kosenok and Sergei (2008), Kos and Messner (2013), and Natha and Sandholm (2019).

$\gamma(s(\mathbf{v}))$, for all $\mathbf{v} \in V$, and since the individual rationality and budget balance conditions hold for any outcome realization, such μ also satisfies (2). The proof for the case of *ex post* Nash equilibrium or Bayes Nash equilibrium is immediate, see e.g., Myerson (1991). A market may thus indirectly be simply defined by μ. Finally, since the setting here is one of private values *ex post* incentive compatibility of μ implies that each trader has a dominant strategy to behave according to her true reservation value.

References

AMBRUS, A. AND G. EGOROV (2017): "Delegation and Nonmonetary Incentives," *Journal of Economic Theory*, 101–135.

ARROW, K. J. (1963): *Social Choice and Individual Values*, New York, NY: John Wiley and Sons.

AUSUBEL, L. M., P. CRAMTON, AND R. J. DENECKERE (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 3, chap. 50, 1897–1945.

BRANDENBURGER, A. AND E. DEKEL (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189 – 198.

CHATTERJEE, K. AND W. SAMUELSON (1983): "Bargaining under Incomplete Information," *Operations Research*, 31, 835–851.

CRAMTON, P., R. GIBBONS, AND P. KLEMPERER (1987): "Dissolving a Partnership Efficiently," *Econometrica*, 55, 615–632.

DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): "The implementation of social choice rules: Some general results on incentive compatibility," *The Review of Economic Studies*, 46, 185–216.

D'ASPREMONT, C. AND L.-A. GERARD-VARET (1979): "Incentives and incomplete information," *Journal of Public Economics*, 11, 25–45.

EECKHOUT, J. AND P. KIRCHER (2018): "Assortative Matching with Large Firms," *Econometrica*, 86.

GIBBARD, A. (1973): "Manipulation of Voting Schemes: A General Result," *Econometrica*, 41, 487–501.

GREEN, J. AND J.-J. LAFFONT (1979): *Incentives in Public Decision Making*, North-Holland, Amsterdam.

HAGERTY, K. M. AND W. P. ROGERSON (1987): "Robust Trading Mechanisms," *Journal of Economic Theory*, 42, pp. 94–107.

HOLMSTROM, B. (1977): "On incentives and control in organizations Ph. D," Ph.D. thesis, thesis, Stanford University.

HOLMSTRÖM, B. M. AND R. B. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, pp. 1799–1819.

KOS, N. AND M. MESSNER (2013): "Extremal incentive compatible transfers," *Journal of Economic Theory*, 134–164.

KOSENOK, G. AND S. SERGEI (2008): "Individually rational, budget-balanced mechanisms and allocation of surplus," *Journal of Economic Theory*, 140, 126–161.

LEDYARD, J. O. (1978): "Incentive Compatibility and Incomplete Information," *Journal of Economic Theory*, 18, pp. 171–189.

MCAFEE, R. P. (1992): "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," *Journal of Economic Theory*, 56, 266–293.

MENZIO, G. AND N. TRACHTER (2015): "Equilibrium Price Dispersion with Sequential Search," *Journal of Economic Theory*, 160, 188 – 215.

MERTENS, J.-F. AND S. ZAMIR (1985): "Formulation of Bayesian analysis for games with incomplete information," *International Journal of Game Theory*, 14, 1–29.

MYERSON, R. B. (1979): "Incentive compatibility and the bargaining problem," *Econometrica: journal of the Econometric Society*, 61–73.

——— (1991): *Game Theory*, Harvard University Press.

MYERSON, R. B. AND M. A. SATTERTHWAITE (1983): "Efficient Mechanism for Bilateral Trading," *Journal of Economic Theory*, 29, pp. 265–281.

NATHA, S. AND T. SANDHOLM (2019): "Efficiency and budget balance in general quasi-linear domains?" *Games and Economic Behavior*, 113, 673–693.

SATTERTHWAITE, M. A. (1975): "Strategy-Proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions," *Journal of Economic Theory*, 10, 187–217.

STIEGLER, G. J. (1971): "The Theory of Economic Regulation," *The Bell Journal of Economics and Management Science*, 2, 3–21.

Čopič, J. (2019a): "Constrained Efficiency," *Mimeo*.

——— (2019b): "Scarce information and constrained efficiency," *Mimeo*.

Čopič, J., M. Castellani, and v. Velkavrh (2019): "Brokers and bid-ask spread," *Mimeo*.

Čopič, J. and C. Ponsatí Obiols (2016): "Optimal robust bilateral trade: risk neutrality." *Journal of Economic Theory*, 163, 276–287.