

Efficient Bilateral Trade*

Rod Garratt and Marek Pycia

March 2016

Abstract

We re-examine the canonical question of Myerson and Satterthwaite (1983) whether two parties can trade an indivisible good in a Pareto efficient way when they are both privately-informed about their valuations for the good. Relaxing the assumption that utilities are quasi-linear, we show that efficient trade is generically possible if agents' utility functions are not too responsive to private information. In natural examples efficient trade is possible even when agents' utility functions are highly responsive to their private information. The analysis relies on new methods we introduce.

1 Introduction

Can a seller and a buyer of an indivisible good trade efficiently if they are privately informed about their valuations for the good and if ex ante either of them might have the higher valuation? The theorem of Myerson and

*The views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of New York or the Federal Reserve System. We are still tweaking the exposition. For their comments, we would like to thank Andrew Atkeson, Larry Ausubel, Simon Board, Peter Cramton, Erik Eyster, Srihari Govindan, Peter Hammond, Hugo Hopenhayn, Philippe Jehiel, Igor Kopylov, Moritz Meyer-ter-Vehn, Ichiro Obara, Gregory Pavlov, Jim Peck, John Riley, Ronny Razin, Stergios Skaperdas, Tayfun Sonmez, Guofu Tan, Tom Wilkenning, Steven Williams, William Zame, and the audiences at SWET 2014, LSE, UCL, NBER Decentralization Conference 2014, SING 2014, UCLA, Rochester, BC, the Bank of England, Amherst, NEGTE 2014, Rutgers, LA Theory Conference 2014, UC Irvine, U of Maryland and U of Bayreuth. Garratt: FRBNY, Pycia: UCLA.

Satterthwaite (1983), a central result of the theory of mechanism design, provides a negative answer to this question. Assuming that agents have quasi-linear utility functions, Myerson and Satterthwaite showed that no Bayesian incentive-compatible, individually rational, non-subsidized mechanism is ex-post Pareto efficient. The reason is that in an incentive compatible mechanism, each agent needs to be provided with rents commensurate with her type, and the gains from efficient trade are not sufficient to cover the rents of both parties. Their impossibility theorem had a large impact on the economics literature and the practice of market design, and it offers a stark contrast to Coase's (1960) claim that markets lead to efficient outcomes when property rights are unambiguously established and there are no transaction costs.¹

We examine the possibility of efficient trade allowing for risk aversion and wealth effects, phenomena which are assumed away in the standard quasilinear analysis and which nonetheless play an important role in bilateral trade of valuable assets. Surprisingly, we establish that the presence of risk aversion or wealth effects greatly improves the possibilities for efficient trade.²

Our main result establishes that efficient trade is generically possible as long as either agents' utilities are not too dependent on private information, or the asymmetry of information is not too large. We provide a Bayesian incentive-compatible and interim individually rational mechanism that, under these conditions, is ex-post Pareto efficient and that does not rely on any subsidies or on budget breaking by third parties.³ The restriction on private information is necessary in this general possibility result as the scope

¹See, for instance, Milgrom's (2004) discussion of the role the impossibility theorem played in the FCC deliberations on the first US spectrum auctions and Loertscher, Marx, and Wilkening's (2013) discussion of how the impossibility theorem led to the focus of market design on primary markets. Cramton, Gibbons, and Klemperer (1987) contains the original research that provided the argument to the FCC against using lotteries.

²While we describe the problem in terms of trade of an indivisible good, our results remain true for trading a divisible good, or for trading multiple goods against each other. The tools we develop are useful beyond bilateral trade models.

³Like Myerson and Satterthwaite we look at ex post efficiency in the sense that we evaluate efficiency assuming that we know the agents' types. With respect to the resolution of the lotteries, our contracts are efficient not only ex post but also ex ante.

of efficient trade diminishes in the limit as we approach the quasi-linear model. We complement this general possibility result by showing that when risk-aversion or wealth effects are substantive, or when the good traded is valuable, then efficient trade may be possible even if agents' utilities are highly dependent on their private information. In particular, our results imply that the central impossibility insight of mechanism design hinges on the assumption of quasi-linear utilities.⁴

Efficient trade remains possible if agents' private information determines not only their valuations for the good but also the marginal utility of money. Our main result remains valid in this more general setting as long as (i) the traded good is normal (a good is normal if each player's reservation price for the good increases with the agent's wealth; see, for instance, Cook and Graham 1977), and (ii) the semi-elasticity of marginal utility of money with respect to private information is constant and dominated by the semi-elasticity of marginal utility of the good with respect to private information.⁵ For instance, agents endowed with Cobb-Douglas utilities can trade efficiently regardless of the extent of informational asymmetry. Our analysis furthermore implies that efficient trade might be possible even if neither of the assumptions (i) and (ii) is satisfied.⁶

⁴We consider environments where the item that is being allocated is wholly owned by one agent, whom we refer to as the seller. Within the quasilinear framework, Cramton, Gibbons, and Klemperer (1987) show that the buyer-seller ownership structure is also essential to the impossibility result. When the initial ownership is not too asymmetric, full efficiency is possible. The source of inefficiency is the combination of asymmetric information and asymmetric ownership. Even a modest amount of joint ownership is enough to restore the possibility of efficient trade. For example, in the setting of a divisible good, if the initial quantities are not too asymmetric, efficient trade is possible.

⁵The normality condition for an indivisible good is a natural counterpart of normality for divisible commodities. Cook and Graham require that each player's reservation price for the good strictly increases with the agent's wealth. Thus their condition does not hold under quasi-linear utility. Our results however continue to hold for generic utility profiles if we relax the Graham and Cook normality condition to require only that each player's reservation price for the good weakly increases with the agent's wealth. Such a relaxed normality condition is satisfied by quasi-linear preferences.

⁶In particular, the trading mechanism we construct is efficient, individually rational, budget-balanced, and makes truthful reporting a solution to the first-order condition of the agents' maximization problem whether or not assumption (ii) is satisfied. Assumption (ii) ensures that the second-order condition of the agents' maximization problem is satisfied for all possible distributions of private information.

Why is efficient trade possible when agents are risk averse, but it cannot be attained with quasi-linear preferences? There are two parts to the explanation. First, risk aversion and wealth effects open up an additional source of efficiency gains. Suppose the seller's utility function over money and the object is the same as the buyer's and suppose that seller's and buyer's money holdings are such that they derive the same utility from their endowments. Consider giving the good to the buyer with a small probability, while compensating the seller with a money transfer in the state when the seller keeps the good. Normality means that the money transfer needed to compensate the seller for the small probability of giving up the good is less than the money transfer the buyer is willing to make in return for the same probability of obtaining the good. Hence, such a lottery contract is Pareto improving. The role for non-degenerate lotteries does not emerge in the case of quasi-linear preferences where efficiency means assigning the item to the highest-value agent with probability one.

Second, the dependence of efficient trades on lotteries introduces countervailing incentives to misrepresenting types that are absent in the quasilinear setting. In the quasilinear setting there are no countervailing incentives to misrepresenting types. If we start with the ex post efficient allocation in the Myerson and Satterthwaite (1983) environment with people telling the truth, it cannot be incentive compatible. If the seller overstates her value by a little bit, it results in a higher price when she sells, but the accompanying reduction in the probability of sale only matters in the case where she is almost indifferent to selling. In our setting, if the seller misreports her type there is a real tradeoff.⁷

To establish our possibility result, we need to develop a methodology to study Pareto efficient trading mechanisms in settings with risk-averse agents. For instance, in symmetric settings, Pareto efficiency requires randomization when agents' valuations for the good are nearly equal; thus the mechanism design problem we study is very different from the one studied by Myerson and Satterthwaite precisely in the range of types that underlies their impos-

⁷We thank Jim Peck for suggesting that we frame the explanation in terms of countervailing incentives.

sibility result. Unlike in their quasilinear setting, in ours the size of money transfers conditional on allocation are then uniquely determined by efficiency but who obtains the allocation of the good itself is not. The key to constructing the Pareto efficient mechanism is a judicious choice of the probability of allocating the good to each of the trading agents; in a direct mechanism the probability needs to respond to agents' reports in a way that ensures that truthful reporting is Bayesian incentive compatibility and interim individually rational. We reduce the problem of constructing such a probability function to a non-standard system of partial differential equations and offer a constructive way to solve this system of equations.⁸

The question when efficient trade is possible is important and has been extensively studied. Two crucial assumptions have been recognized already by Myerson and Satterthwaite: their result requires that the distribution of types is continuous and that buyer's value is always higher than the seller's value.⁹ Gresik and Satterthwaite (1983) and Williams (1999) extend the impossibility result to symmetric settings with many agents while Makowski and Mezzetti (1994) show that for an open set of asymmetric distributions trade can be efficient provided there are multiple buyers.¹⁰ Wilson (1985), Makowski and Ostroy (1989), Rustichini, Satterthwaite, and Williams (1994), Reny and Perry (2006), Cripps and Swinkels (2006), and many others establish that trade is asymptotically efficient as the number of buyers and sellers becomes large. McAfee (1991) shows that efficient trade can be possible when the ex ante gains from trade are large and the two trading parties have access to an uninformed budget breaking third agent.¹¹

At the same time, we know of no successful attempt to go beyond My-

⁸Each of the partial differential equations needs to be satisfied only on average. In this sense, this system of equations resembles the problem of finding an allocation rule in reduced form auctions (though the latter problem involves no differentiation); see e.g. Border (1991) and Che, Kim, and Mierendorff (2013).

⁹The case of discrete distributions is studied by Matsuo (1989) and Kos and Manea (2008).

¹⁰This observation has been extended to settings with multiple buyers and sellers by Williams (1999) and Schweizer (2006).

¹¹McAfee (1991) studied the problem of trading divisible goods; see (Riley 2012) for an analysis with indivisible goods.

erson and Satterthwaite (1983) and demonstrate the possibility of fully efficient mechanisms in their original context of two expected-utility maximizing agents, one buyer and one seller. What was demonstrated is the possibility of approximate efficiency in two contexts. Chatterjee and Samuelson (1983) show that double-auctions are asymptotically efficient in the limit as the agents become infinitely risk-averse.¹² And, McAfee and Reny (1992) show that when private values are correlated (and a hazard rate assumption is satisfied) a judicious use of Cremer and McLean (1988) lotteries allows the parties to reduce their incentives to misreport so as to permit outcomes as close to efficiency as desired.¹³

While ours is likely the first paper to study efficient Bayesian incentive compatible bilateral trade mechanisms for an indivisible normal good, it follows a rich mechanism design literature studying risk aversion. This literature generally assumed that the seller is uninformed (e.g. Garratt (1999) and Baisa (2013)), and it often restricted attention to deterministic mechanisms (e.g. Holt (1980) and Morimoto and Serizawa (2014)), or focused on revenue-maximization (e.g. Matthews (1983), Maskin and Riley (1984)). The closest prior paper is Garratt (1999) who studies the Pareto frontier and shows that random mechanisms can dominate deterministic ones in a complete information setting; in contrast, we allow for private information.

The rest of the paper is organized as follows. In Section 2 we present our model and assumptions. Section 3 contains our main results. Sup-

¹²Chatterjee and Samuelson (1983) also show that double auction is not efficient, and Chatterjee (1982) shows that no mechanism in a large class of incentive compatible mechanisms is efficient; these two papers are important precursors of Myerson and Satterthwaite's impossibility theorem. After Myerson and Satterthwaite, many authors provided alternative proofs of their impossibility result, see for instance Williams (1999) and Krishna and Perry (1998). Recent work looks also at bilateral trade among agents who are not expected utility maximizers: in a paper simultaneous and independent from ours, Wolitzky (2015) analyzes bilateral trade among max-min agents, and, following our work, Crawford (2016) analyzes bilateral trade between k -level agents.

¹³For simplicity, we formulate our main theorem in the context of risk averse agents, but we allow any level of risk aversion, including arbitrarily small risk aversion. Furthermore, as demonstrated by the example of Section 5.1, the underlying insight does not rely on risk aversion. Also, our results allow for both independence and correlation, and they do not rely on large bets in the spirit of Cremer and McLean: all lotteries we employ are bounded by agents' wealth levels.

pose that the traded object is normal and that agents' types impact their marginal utility of the good more than the marginal utility of money, and that either the utility functions are not too dependent on types, or that the supports of agents' types are not too large. Then, for any initial distribution of money holdings, except possibly one, we construct a mechanism that is incentive-compatible, individually rational and efficient. Our main result implies that efficient trade is possible in problems arbitrarily close to Myerson and Satterthwaite's quasilinear setting. In Section 4, we consider the case of separable utility and compute and example for the case where the utility over money holdings is characterized by the log function. In Section 5 we discuss the generic impossibility of ex post implementation and show, by an example, that in some settings efficient trade is possible for any distribution of agents' types, and for any distribution of agents' initial wealth levels. In the example, efficient trade is ex post implementable and the mechanism works regardless of whether or not endowed money holdings are private information. Concluding remarks are found in Section 6.

2 Model and Assumptions

A seller is endowed with a good and money endowment m_s while the buyer has money endowment m_b . We denote aggregate money holdings by $M = m_s + m_b$. We assume the seller's type (cost) $c \in [\underline{c}, \bar{c}]$ and the buyer's type (value) $v \in [\underline{v}, \bar{v}]$. Furthermore, we assume the utility function $u(x, m, \theta)$ of each agent is monotonic in having the indivisible good (x) and in money (m), strictly concave in money, and twice continuously differentiable in money and in type.¹⁴ For convenience, we extend the utility function notation to lotteries over the good: $u(x, m, \theta) = xu(1, m, \theta) + (1 - x)u(0, m, \theta)$ for $x \in [0, 1]$.

We assume that the good is **normal** in the standard sense (see Cook and Graham 1977): for any type θ and any relevant money levels $m, p, \epsilon > 0$, if $u(0, m, \theta) = u(1, m - p, \theta)$ then $u(0, m + \epsilon - p, \theta) < u(1, m + \epsilon, \theta)$. Normality

¹⁴As usual, the utility of money derives from other goods the agent might purchase.

captures the intuition that the more money an agent has, the more she is willing to pay for the good.

We also impose the following assumption on how agents' utilities respond to their types. We assume that $\frac{\partial}{\partial \theta} \log \left(\frac{\partial}{\partial m} u(x, m, \theta) \right)$ does not depend on m , and that for any $x \in [0, 1]$, $m \in [0, M]$, and any type θ , we have

$$\frac{\partial}{\partial \theta} \log \left(\frac{\partial}{\partial x} u(x, m, \theta) \right) > \frac{\partial}{\partial \theta} \log \left(\frac{\partial}{\partial m} u(x, m, \theta) \right) \quad (1)$$

That is we want the type-elasticity of the marginal value of the good to exceed the type-elasticity of the marginal utility of money, and the latter elasticity to be constant in money. Both components of this assumption are automatically satisfied when utilities are quasilinear in money, and higher types have higher utility from consuming the good. The assumption is also satisfied when agents' utilities are additively separable, $u(x, m, \theta) = \theta x + v(m)$. Furthermore, this assumption is only needed for our analysis of the second-order condition of the mechanism we construct, where it is sufficient but not necessary. The necessary condition is more complex, and is provided in the analysis of the second-order condition.

The normality assumption allows us to determine some features of the Pareto frontier (See Figure 1).¹⁵ To describe the frontier, let us fix agents' types c and v . The frontier is the upper envelope of the curves

$$\mathcal{C}_S = \{(u(1, m, c), u(0, M - m, v)) : m \in [0, M]\}$$

and

$$\mathcal{C}_B = \{(u(0, m, c), u(1, M - m, v)) : m \in [0, M]\}.$$

The curve \mathcal{C}_S traces the utilities when the seller has the good, while the curve \mathcal{C}_B traces the utilities when the buyer has the good. Since we assume that it is better to have the good than not to have it, the curve \mathcal{C}_S starts higher than \mathcal{C}_B on the axis of seller's utilities (the vertical axis), and \mathcal{C}_S ends lower than \mathcal{C}_B on the axis of buyer's utilities (the horizontal axis). We also

¹⁵See Garratt (1999) for a more detailed discussion of the Pareto frontier for normal goods.

assume that each agent prefers to have all of the money and no good to the good and no money: $u_s(0, M, c) > u_s(1, 0, c)$ and $u_b(0, M, v) > u_b(1, 0, v)$.¹⁶ This assumption ensures that the two curves intersect, and the normality assumption implies that \mathcal{C}_S intersects \mathcal{C}_B only one time from above.¹⁷




Figure 1: Pareto Frontier with a Normal Good.

As we move along the Pareto frontier from the seller's most preferred point to the buyer's most preferred point, we start on the curve \mathcal{C}_S and we end on the curve \mathcal{C}_B . The point at which the curves \mathcal{C}_S and \mathcal{C}_B intersect cannot be part of the frontier because normality implies that at this point \mathcal{C}_S intersects \mathcal{C}_B strictly from above, and hence a randomization over any point just to the left of the intersection and any point just to the right of the intersection is strictly preferred to the intersection point by both trading parties. The frontier thus contains a flat part consisting of randomizations between two points: $S(c, v) \in \mathcal{C}_S$ and $B(c, v) \in \mathcal{C}_B$, where the seller strictly prefers S to B while the buyer strictly prefers B to S . The strict concavity of u in money implies that these two points are uniquely determined. We call them the critical points. Let $m^S(c, v)$ denote the seller's money holdings at S

¹⁶This is not necessary for the analysis, but it simplifies the presentation and validates the cases depicted in the figures. Note that in the figures the origin is not necessarily the point $(0,0)$. It is the point $(u_b(1, 0, v), u_s(1, 0, c))$.

¹⁷The uniqueness of the intersection point is the main implication of normality in the paper. Without normality the two curves could intersect multiple times.

and $m^B(c, v)$ denote the buyer's money holdings at B . The critical levels of money holdings depend upon agents' reports of their types. The dependence of the points $B(c, v)$ and $S(c, v)$ on (c, v) is continuously differentiable by the assumed regularity of u and its strict concavity in money. We assume that these two points are either at the boundary (that is involve money levels of M and 0) or they are both internal. When the critical points S and B are internal, the money levels $m^S(c, v)$ and $m^B(c, v)$ are uniquely determined by the following equations

$$\frac{\frac{\partial}{\partial m} u(1, m^S, c)}{\frac{\partial}{\partial m} u(0, M - m^S, v)} = \frac{\frac{\partial}{\partial m} u(0, M - m^B, c)}{\frac{\partial}{\partial m} u(1, m^B, v)} = \frac{u(1, m^S, c) - u(0, M - m^B, c)}{u(1, m^B, v) - u(0, M - m^S, v)}. \quad (2)$$

These equations express the fact that the randomization interval is tangent to the Pareto frontier at both critical points S and B .

In the appendix we show that in the internal case our assumptions imply that

$$\frac{\partial}{\partial \theta} m^{-\theta}(c, v) > 0 > \frac{\partial}{\partial \theta} m^{\theta}(c, v), \quad (3)$$

(in the corner case the two inequalities become equalities). Thus, agents' money holdings at the critical point of the Pareto frontier the agent prefers are decreasing in agent's type, while the agent's money holdings at the critical point the other agent prefers are increasing in agent's own type.

3 Possibility Theorems on Efficient Trade

We now show that efficient trade is possible for the class of utility functions described in Section 2. We formulate our possibility result in three related ways. First, efficient trade is possible if agents' utilities are not too dependent on their types.

Theorem 1. *Fix $c^*, v^* \in [0, 1]$ and $u^s(\cdot, \cdot; c^*)$ and $u^b(\cdot, \cdot; v^*)$. For every profile of money endowments but one, there is $\delta > 0$ such that if*

$$\max_{\theta \in [0, 1], m \in [0, M], x \in \{0, 1\}} |u(x, m, \theta) - u(x, m, \theta^*)| < \delta, \quad (4)$$

then there is an incentive-compatible, individually-rational, and budget-balanced mechanism that generates efficient trade.

As an immediate corollary we obtain

Corollary 1. *Fix $(c^*, v^*) \in (0, 1)^2$ and function u . For any profile of money endowments but one, there are intervals $(\underline{c}, \bar{c}) \ni c^*$ and $(\underline{v}, \bar{v}) \ni v^*$ such that: if agents draw their types from arbitrary distributions on $(\underline{c}, \bar{c}) \times (\underline{v}, \bar{v})$, then there is an incentive-compatible mechanism, individually-rational, and budget-balanced mechanism that generates efficient trade.*

This corollary obtains because by taking the intervals $(\underline{c}, \bar{c}) \ni c^*$ and $(\underline{v}, \bar{v}) \ni v^*$ to be sufficiently small, and re-scaling them to $[0, 1]$, we can ensure condition (4).

A special case of interest obtains when $c^* = v^*$. In this case, the corollary takes the following simpler form:

Corollary 2. *Fix any $\theta^* \in (0, 1)$ and function u . For any profile of money endowments but one, there is an interval $(\underline{\theta}, \bar{\theta}) \ni \theta^*$ such that for any distribution of agents' types on $(\underline{\theta}, \bar{\theta}) \times (\underline{\theta}, \bar{\theta})$, there is an incentive-compatible, individually-rational, and budget-balanced mechanism that generates efficient trade.*

In fact, for a given utility function u the range of types that permit efficient trade can be quite large, as is illustrated by Example 1 in Section 6.

3.1 Proof of Theorem 1

Fix a distribution of endowed money holdings other than $(m^S(c^*, v^*), M - m^S(c^*, v^*))$. There are three cases depending on how much money the seller initially has.

In the first case, the seller's endowed money holdings are strictly above $m^S(c^*, v^*)$, then it remains so for utilities close to u^s, u^b , and hence there is $\delta > 0$ that guarantees that the no-trade mechanism is efficient and satisfies all our other requirements.

In the second case, the seller's endowment is such that the seller's utility is strictly below $u(0, M - m^B(c^*, v^*), c^*)$. Then there is a point on the

Pareto frontier of (c^*, v^*) that strictly dominates the agents' utility at the initial endowments (point F in Figure 2). At this point on the frontier the seller has no good and has money holdings $m^S + t$ for some constant transfer t , while the buyer has the good and money holdings $m^B - t$. In particular, the pre-trade seller's utility is strictly below $u(0, m^S + t, c^*)$ and the pre-trade buyer's utility is strictly below $u(1, m^B - t, v^*)$. These bounds on agents' pre-trade utility remain true for type profiles close to (c^*, v^*) . There is then $\delta > 0$ that guarantees that the mechanism that allocates the good and money $m^B - t$ to the buyer, and money $m^S + t$ (without good) to the seller is Pareto efficient, individually rational, and does not require a subsidy. Furthermore, this mechanism is incentive compatible as it does not rely on agents reports.¹⁸




Figure 2: Individually Rational Part of the Pareto Frontier. The Case of Trade at Fixed Price.

In the third case, the seller's endowment is intermediate, that is the seller's money holdings are strictly below $m^S(c^*, v^*)$ but seller's utility is weakly above $u(0, M - m^B(c^*, v^*), c^*)$. This is the main case, and the

¹⁸ A similar fixed-terms-of-trade mechanism delivers efficient trade whenever there exists a point on the Pareto frontier that is strictly preferred by both the buyer and seller to status quo, and strictly more favorable to the buyer than having the good and money $m^B(c^*, v^*)$.

remainder of the proof is devoted to its analysis.

In this third case, there is a point $F = F(c^*, v^*)$ on the flat part of the frontier strictly between $S(c^*, v^*)$ and $B(c^*, v^*)$ (see Figure 3) that is strictly preferred by the buyer and the seller to the initial situation, $(u_s(1, m_s; c^*), u_b(0, m_b; v^*))$. The point F is determined by $\pi(c^*, v^*) = \pi^* \in (0, 1)$ as follows: the utility pair F corresponds to the seller having the good and wealth $m^S(c^*, v^*)$ with probability π^* , and the buyer having the good and wealth $m^B(c^*, v^*)$ with probability $1 - \pi^*$.




Figure 3: Individually Rational Part of the Pareto Frontier. The Case of Trade at Varying Prices.

For small δ , the critical money holdings $m^S(c, v)$ and $m^B(c, v)$ are nearby $m^S(c^*, v^*)$ and $m^B(c^*, v^*)$, and there are $\pi(c, v)$ nearby π^* such that the corresponding $F(c, v)$ is strictly preferred by the agents to the initial situation. If both $m^S(c, v)$ and $m^B(c, v)$ are locally constant around (c^*, v^*) then a mechanism with fixed $\pi(c, v) = \pi^*$ satisfies our postulates. Let us thus assume that at least one $m^S(c, v)$ and $m^B(c, v)$ varies in c, v . By (3) and its qualifying discussion, one of the two functions have non-zero partials throughout the domain.

The crux of the remainder of the argument is to show that $\pi(c, v)$ can be defined in such a way that in a Bayesian Nash equilibrium of the budget-

balanced mechanism that assigns the seller the good and money holdings $m^S(c, v)$ with probability $\pi(c, v)$ and assigns the good and money $m^B(c, v)$ to the buyer with probability $1 - \pi(c, v)$, both agents report their true types. (Note that in this mechanism the buyer gets money $M - m^S(c, v)$ when the seller gets the good, and the seller gets money $M - m^B(c, v)$ when the buyer gets the good).

We thus need to find a function π such that for the seller

$$\Pi^S(c, \hat{c}) = E_v \left(\pi(\hat{c}, v) u(1, m^S(\hat{c}, v), c) + (1 - \pi(\hat{c}, v)) u(0, M - m^B(\hat{c}, v), c) \right)$$

is maximized at $\hat{c} = c$, and similarly for the buyer,

$$\Pi^B(v, \hat{v}) = E_c \left(\pi(c, \hat{v}) u(0, M - m^S(c, \hat{v}), v) + (1 - \pi(c, \hat{v})) u(1, m^B(c, \hat{v}), v) \right)$$

is maximized at $\hat{v} = v$. In order to guarantee this we will construct π such that the first order condition is satisfied for truthful reporting, and the second order condition is satisfied at all points satisfying the first order condition. We will take $\pi(c^*, v^*)$ to be $\pi^* \in (0, 1)$ given above, thus guaranteeing that $\pi(c, v) \in (0, 1)$ for small $\delta > 0$, and that the individual rationality is satisfied for small $\delta > 0$. As an aside, let us note that the individual rationality and the requirement that π is a probability constraints our mechanism to be well-behaved only locally around c^*, v^* ; the incentive compatibility conditions could be made to be globally satisfied.

3.1.1 The First Order Condition

Assuming truthful reporting by the other agent, the first order condition for the seller is

$$\begin{aligned}
0 = E_v \Bigg\{ & \left[\frac{\partial}{\partial \hat{c}} \pi(\hat{c}, v) \right] u(1, m^S(\hat{c}, v), c) \\
& + \pi(\hat{c}, v) \left[\frac{\partial}{\partial m} u(1, m^S(\hat{c}, v), c) \right] \left[\frac{\partial}{\partial \hat{c}} m^S(\hat{c}, v) \right] \\
& - \left[\frac{\partial}{\partial \hat{c}} \pi(\hat{c}, v) \right] u(0, M - m^B(\hat{c}, v), c) \\
& - (1 - \pi(\hat{c}, v)) \left[\frac{\partial}{\partial m} u(0, M - m^B(\hat{c}, v), c) \right] \left[\frac{\partial}{\partial \hat{c}} m^B(\hat{c}, v) \right] \Bigg\}, \quad (5)
\end{aligned}$$

and the first order condition for the buyer is

$$\begin{aligned}
0 = E_c \Bigg\{ & \left[\frac{\partial}{\partial \hat{v}} \pi(c, \hat{v}) \right] u(0, M - m^S(c, \hat{v}), v) \\
& - \pi(c, \hat{v}) \left[\frac{\partial}{\partial m} u(0, M - m^S(c, \hat{v}), v) \right] \left[\frac{\partial}{\partial \hat{v}} m^S(c, \hat{v}) \right] \\
& - \left[\frac{\partial}{\partial \hat{v}} \pi(c, \hat{v}) \right] u(1, m^B(c, \hat{v}), v) \\
& + (1 - \pi(c, \hat{v})) \left[\frac{\partial}{\partial m} u(1, m^B(c, \hat{v}), v) \right] \left[\frac{\partial}{\partial \hat{v}} m^B(c, \hat{v}) \right] \Bigg\} \quad (6)
\end{aligned}$$

We want $\hat{c} = c$ to satisfy the seller's first order condition and $\hat{v} = v$ to satisfy the buyer's first order condition, and hence the two conditions give us a system of PDE equations on $\pi(c, v)$.¹⁹ These equations take the form

$$\begin{aligned}
E_v \left[S_1(c, v) \frac{\partial}{\partial c} \pi(c, v) + S_2(c, v) \pi(c, v) \right] &= \phi(c), \\
E_c \left[B_1(c, v) \frac{\partial}{\partial c} \pi(c, v) + B_2(c, v) \pi(c, v) \right] &= \psi(v),
\end{aligned}$$

¹⁹To ensure that the coefficient in front of $\frac{\partial}{\partial c} \pi(c, v)$ is positive, we multiply the second equation by (-1) before calculating B_1 , B_2 , and ψ below.

where the coefficients in front of $\frac{\partial}{\partial c}\pi$ and $\frac{\partial}{\partial v}\pi$ are

$$\begin{aligned} S_1(c, v) &= u(1, m^S(c, v), c) - u(0, M - m^B(c, v), c) > 0, \\ B_1(c, v) &= u(1, m^B(c, v), v) - u(0, M - m^S(c, v), v) > 0, \end{aligned}$$

the coefficients in front of π are

$$\begin{aligned} S_2(c, v) &= \left[\frac{\partial}{\partial w} u(1, m^S(c, v), c) \right] \left[\frac{\partial}{\partial c} m^S(c, v) \right] + \left[\frac{\partial}{\partial w} u(0, M - m^B(c, v), c) \right] \left[\frac{\partial}{\partial c} m^B(c, v) \right], \\ B_2(c, v) &= \left[\frac{\partial}{\partial w} u(1, m^B(c, v), v) \right] \left[\frac{\partial}{\partial v} m^B(c, v) \right] + \left[\frac{\partial}{\partial w} u(0, M - m^S(c, v), v) \right] \left[\frac{\partial}{\partial v} m^S(c, v) \right], \end{aligned}$$

and the functions ϕ, ψ are given by

$$\begin{aligned} \phi(c) &= E_v \left\{ \left[\frac{\partial}{\partial w} u(0, M - m^B(c, v), c) \right] \left[\frac{\partial}{\partial c} m^B(c, v) \right] \right\}, \\ \psi(v) &= -E_c \left\{ \left[\frac{\partial}{\partial w} u(1, m^B(c, v), v) \right] \left[\frac{\partial}{\partial v} m^B(c, v) \right] \right\}. \end{aligned}$$

By assumption u and its derivatives are continuously differentiable. The continuous differentiability of m^S and m^B follows from strict concavity of u and the implicit function theorem (the implicit equations defining m^S and m^B are in the appendix).

The above averaged-out system of PDEs has a solution for any initial condition $\pi(c^*, v^*) = \pi^*$ by the following crucial lemma (proven in the appendix).

Lemma 1. *Let I be a bounded interval of positive length and let F be a joint distribution of (c, v) over domain $I^2 \subseteq \mathbb{R}^2$. Let $S_1(\cdot, \cdot), S_2(\cdot, \cdot)$ and $B_1(\cdot, \cdot), B_2(\cdot, \cdot)$ be functions defined on I^2 , and ϕ, ψ be functions on I . Suppose that all these functions are continuously differentiable and that $S_1, B_1 \neq 0$ for all arguments (c, v) . Then, the system of PDE equations*

$$E_v \left[S_1(c, v) \frac{\partial}{\partial c} \pi(c, v) + S_2(c, v) \pi(c, v) \right] = \phi(c), \quad (7)$$

$$E_c \left[B_1(c, v) \frac{\partial}{\partial v} \pi(c, v) + B_2(c, v) \pi(c, v) \right] = \psi(v), \quad (8)$$

has a solution π for any boundary condition $\pi(c^*, v^*) = \pi^*$. Furthermore, as $\frac{S_2}{S_1}$, $\frac{B_2}{B_1}$, $\frac{\phi}{S_1}$, and $\frac{\psi}{B_1}$ tend to zero, the derivatives of the solution π tend to zero as well.

Finally, notice that because $\pi^* \in (0, 1)$, and because by taking Δ to be sufficiently small we can guarantee that π is near flat, we can find a domain of types such that π takes values in $[0, 1]$ and satisfies the individual rationality conditions on the entire domain.

The above lemma is of independent interest. It tells us that for any marginal distributions of the linear PDE formulas from the lemma, we can find a function that implements these marginal distributions.²⁰

3.1.2 The Second Order Condition

The last thing to check is that the agents objectives satisfy the second-order condition at every point at which the first-order condition is satisfied so that truthful reporting is not only a solution of the first-order condition but also the optimal report. Let us thus check the second-order conditions for the seller; the buyer's problem is analogous.

Since at points at which the first-order condition is satisfied we have

$$0 = \frac{d}{dc} \left(\frac{\partial}{\partial \hat{c}} \Pi^S(c, c) \right) = \frac{\partial}{\partial c} \left(\frac{\partial}{\partial \hat{c}} \Pi^S(c, c) \right) + \frac{\partial}{\partial \hat{c}} \left(\frac{\partial}{\partial \hat{c}} \Pi^S(c, c) \right),$$

the second-order condition for the seller would be implied if we shown that

$$\frac{\partial}{\partial c} \frac{\partial}{\partial \hat{c}} \Pi^S(c, c) > 0.$$

²⁰We have not been able to find this lemma in the literature on partial differential equations. The sufficient conditions for existence of solutions of non-averaged linear PDEs of Thomas (1934) and Mardare (2007) can easily tell us that the lemma is true if $\frac{\partial}{\partial v} \frac{S_2}{S_1} = \frac{\partial}{\partial c} \frac{B_1}{B_2}$, which is satisfied for instance when the coefficients B_i, S_i are all constant, but they are not satisfied in the general case we consider here (which is not surprising as it is much easier to satisfy the PDE equations on average than it is to satisfy them pointwise).

A straightforward calculation shows that $\frac{\partial}{\partial c} \frac{\partial}{\partial \bar{c}} \Pi^S(c, c)$ equals

$$\begin{aligned} & E_v \left\{ \left[\frac{\partial}{\partial c} \pi(c, v) \right] [u_c(1, m^S(c, v), c) - u_c(0, M - m^B(c, v), c)] \right. \\ & + \pi(c, v) \left(\left[\frac{\partial}{\partial m} u_c(1, m^S(c, v), c) \right] \left[\frac{\partial}{\partial c} m^S(c, v) \right] + \left[\frac{\partial}{\partial m} u_c(0, M - m^B(c, v), c) \right] \left[\frac{\partial}{\partial c} m^B(c, v) \right] \right) \\ & \left. - \left[\frac{\partial}{\partial m} u_c(0, M - m^B(c, v), c) \right] \left[\frac{\partial}{\partial c} m^B(c, v) \right] \right\} \end{aligned}$$

We can substitute in for $\frac{\partial}{\partial c} \pi(c, v)$ from the first-order condition obtaining that $\frac{\partial}{\partial c} \frac{\partial}{\partial \bar{c}} \Pi^S(c, c)$ equals $(1 - \pi(c, v)) \frac{\partial}{\partial c} m^B(c, v)$ times

$$\begin{aligned} & \left[\frac{\partial}{\partial m} u(0, M - m^B(c, v), c) \right] [u_c(1, m^S(c, v), c) - u_c(0, M - m^B(c, v), c)] \\ & - \left[\frac{\partial}{\partial m} u_c(0, M - m^B(c, v), c) \right] [u(1, m^S(c, v), c) - u(0, M - m^B(c, v), c)] \end{aligned}$$

minus $\pi(c, v) \frac{\partial}{\partial c} m^S(c, v)$ times

$$\begin{aligned} & \left[\frac{\partial}{\partial m} u(1, m^S(c, v), c) \right] [u_c(1, m^S(c, v), c) - u_c(0, M - m^B(c, v), c)] \\ & - \left[\frac{\partial}{\partial m} u_c(1, m^S(c, v), c) \right] [u(1, m^S(c, v), c) - u(0, M - m^B(c, v), c)] \end{aligned}$$

By assumption we are considering the case when one of the partials $\frac{\partial}{\partial c} m^B(c, v)$, $\frac{\partial}{\partial c} m^S(c, v)$ is non-zero throughout the domain. Thus, (3) implies that the second order condition is satisfied provided both above displayed expressions are strictly positive. Since $m^S \geq M - m^B$, the expressions are positive if

$$\left[\frac{\partial}{\partial m} u(0, m, c) \right] [u_c(1, m', c) - u_c(0, m, c)] - \left[\frac{\partial}{\partial m} u_c(0, m, c) \right] [u(1, m', c) - u(0, m, c)] > 0$$

and

$$\left[\frac{\partial}{\partial m} u(1, m', c) \right] [u_c(1, m', c) - u_c(0, m, c)] - \left[\frac{\partial}{\partial m} u_c(1, m', c) \right] [u(1, m', c) - u(0, m, c)] > 0$$

for all $m' \geq m$. We can re-express the two inequalities as

$$\frac{u_c(1, m', c) - u_c(0, m, c)}{u(1, m', c) - u(0, m, c)} > \frac{\frac{\partial}{\partial m} u_c(0, m, c)}{\frac{\partial}{\partial m} u(0, m, c)}$$

and

$$\frac{u_c(1, m', c) - u_c(0, m, c)}{u(1, m', c) - u(0, m, c)} > \frac{\frac{\partial}{\partial m} u_c(1, m', c)}{\frac{\partial}{\partial m} u(1, m', c)}$$

for all $m' \geq m$. These two inequalities are implied by our assumptions. Let us show it for the first of the two inequalities; the proof of the second one follows the same steps. Let us rewrite the left-hand side as

$$\begin{aligned} \frac{u_c(1, m', c) - u_c(0, m, c)}{u(1, m', c) - u(0, m, c)} &= \frac{[u_c(1, m', c) - u_c(1, m, c)] + [u_c(1, m, c) - u_c(0, m, c)]}{[u(1, m', c) - u(1, m, c)] + [u(1, m, c) - u(0, m, c)]} \\ &= \frac{[u_c(1, m, c) - u_c(0, m, c)] + \int_m^{m'} u_{cm}(0, \tilde{m}, c)}{[u(1, m, c) - u(0, m, c)] + \int_m^{m'} u_m(0, \tilde{m}, c)} \end{aligned}$$

Now, the first inequality of (1) gives

$$\frac{u_c(1, m, c) - u_c(0, m, c)}{u(1, m, c) - u(0, m, c)} > \frac{\frac{\partial}{\partial m} u_c(0, m, c)}{\frac{\partial}{\partial m} u(0, m, c)},$$

and the constancy of $\frac{u_{cm}(0, \tilde{m}, c)}{u_m(0, \tilde{m}, c)}$ in \tilde{m} gives

$$\frac{\frac{\partial}{\partial m} u_c(0, \tilde{m}, c)}{\frac{\partial}{\partial m} u(0, \tilde{m}, c)} = \frac{\frac{\partial}{\partial m} u_c(0, m, c)}{\frac{\partial}{\partial m} u(0, m, c)}.$$

Thus, the left-hand side is a ratio of sums such that the ratio of each summand in the nominator to the corresponding summand in denominator is weakly higher, and in one non-zero measure case strictly higher than the left-hand side above. This ends the proof of Theorem 1.

4 Separable Utility and the Sufficient Departure from Quasi-linearity

As the results from the previous section indicate, any departure from quasi-linear preferences results in efficient trade for sufficiently constrained environments. Corollaries 1 and 2 express these constraints in terms of the range of player types. An alternative approach is to fix the non-preference primitives and ask whether a sufficient departure from quasi-linearity is enough to ensure efficient trade. We demonstrate this possibility for the case where agents have private information about their valuation for the good, but their marginal utility of money is commonly known. More specifically, we assume the agents' utility takes the following separable form $u(x, m; \theta) = \theta x + U(m)$, where U is strictly increasing, concave, and twice differentiable.²¹

Theorem 2. *Fix any continuous joint distribution of types on a compact type space in $(0, \infty)^2$ and any non-equal initial money endowments. There is $\Delta > 0$ such that for all increasing, concave, and twice differentiable functions U such that $|U''|/U' \geq \Delta$, agents with utilities $u(x, m; \theta) = \theta x + U(m)$ can trade efficiently.*

In the case of constant absolute risk aversion (CARA) utility $U(m) = 1 - e^{-\alpha m}$ Theorem 2 says that efficient trade is possible for a sufficiently large degree of risk aversion (i.e., sufficiently large α). For α arbitrarily close to 0, U is arbitrary close to linear (CARA utility is defined to be linear at $\alpha = 0$). Hence the result says that for efficient trade to be guaranteed we need to be sufficiently far away from quasi-linearity. Example 2 in Section 6 demonstrates this result by fixing money endowments and the type distributions and showing that $\alpha = 1$ is sufficient to ensure efficient trade.

²¹For simplicity we focus on the case where U is common to both agents; this restriction is not crucial.

5 Generic Impossibility of Ex post Implementation

So far our analysis has focused on Bayesian implementation. In a related setting in which a seller wants to allocate a normal good to one of a finite number of buyers, and in which, unlike in our setting, the seller has no private information about the good, Baisa (2013) constructs an elegant example of a profile of utility functions such that no strategy-proof mechanism allocates the good in an efficient way. Our analysis of the first order condition of agents' optimization implies the following:

Proposition 1. *When the randomization interval is interior, $m^S, m^B \in (0, M)$, and money endowments are such that efficiency requires randomization, then for generic utility function u no mechanism can implement efficient trade in an ex-post equilibrium.*

The validity of Proposition 1 follows from the fact that there is, in general, no solution to the system of equations (5)-(6) if expectations are not taken. Nevertheless, we can construct a specialized example (Example 3 in Section 6) where the mechanism that achieves efficient trade is strategy proof. This demonstrates that the possibilities for efficient trade in quasilinear settings extend beyond the cases captured by Theorem 1.

6 Examples

Examples 1 and 2 use separable utility of the form $u(x, m; \theta) = \theta x + U(m)$. Equations describing the efficient contract for this case are provided in the Appendix. Before proceeding with the examples, it is worth highlighting a special feature of the separable case: in the optimal contract each player's money allocation is invariant across states. So each player receives the same money allocation regardless of whether or not they consume the indivisible good. This fact allows an interesting interpretation of the contracts specified in Examples 1 and 2 as examples of a weighted Vickrey-Clarke-Groves mechanism.²² Define the efficient allocation to be that which maximizes the

²²This interpretation was suggested to us by Gregory Pavlov.

weighted sum of utilities where weights are the reciprocals of the agent's private values. Then utility is quasilinear in the probability of receiving the good and VCG transfers, expressed in terms of probability units, can be specified that ensure truth telling and are budget balanced in the sense that probabilities sum to 1. Of course, these "transfers" leave each agent with exactly the consumption probabilities determined by our mechanism.

Example 1: Log preferences

If $U(m) = \log(m)$, then equations (18) and (19) yield $m^B(c, v) = \frac{Mc}{c+v}$ and $m^S(c, v) = \frac{Mv}{c+v}$. Moreover, $\frac{\partial}{\partial c} U(m^S(c, v)) = -\frac{1}{(c+v)}$ and $\frac{\partial}{\partial v} U(m^B(c, v)) = -\frac{1}{(c+v)}$. Substituting these expressions into (20) and (21) yields

$$0 = E_v \left[\frac{\partial \pi(c, v)}{\partial c} c - \frac{1}{c+v} \right] \quad (9)$$

$$0 = E_c \left[-\frac{\partial \pi(c, v)}{\partial v} v - \frac{1}{c+v} \right] \quad (10)$$

We now use the proof of Lemma 1 to compute the solution. Suppose that c, v are distributed independently and uniformly on $[2, 100]$. Define

$$\psi(v) = -E_c \left[-\frac{1}{c+v} \right] = \int_2^{100} \frac{1}{c+v} \frac{1}{98} dc = \frac{\log(100+v) - \log(2+v)}{98}$$

and

$$\phi(c) = E_v \left[-\frac{1}{c+v} \right] = -\int_2^{100} \frac{1}{c+v} \frac{1}{98} dv = \frac{-\log(100+c) + \log(2+c)}{98}.$$

Following the construction in the proof of Lemma 1, we can set

$$\Delta^s(c, v) = 1, \quad \Delta^b(c, v) = 1.$$

The function $b(\cdot)$ is now given by the ODE

$$\begin{aligned}\psi(v) &= E_c \left[B_1(c, v) \Delta^b(c, v) \right] b'(v) + E_c \left[B_1(c, v) \frac{\partial}{\partial v} \Delta^b(c, v) + B_2(c, v) \Delta^b(c, v) \right] b(v) \\ &= E_c[v] b'(v) = vb'(v),\end{aligned}$$

and thus

$$b'(v) = \frac{\log(100 + v) - \log(2 + v)}{98v}.$$

Similarly, $\phi(c) = cs'(c)$, and thus

$$s'(c) = \frac{-\log(100 + c) + \log(2 + c)}{98c}.$$

We use the initial condition $\pi^* = .5$. Then, the seller retains the item with probability .5 when the reports are 51, 51. Variations in the seller's report adjusts the probability by the amount

$$\int_c^{51} s'(x) dx = \frac{1}{98} \int_c^{51} \frac{-\log(100 + x) + \log(2 + x)}{x} dx$$

and variations in the buyer's report adjusts the probability by the amount

$$\int_v^{51} b'(x) dx = \frac{1}{98} \int_v^{51} \frac{\log(100 + x) - \log(2 + x)}{x} dx.$$

So the probability that the seller gets the item if the seller reports c and the buyer reports v is

$$\pi(c, v) = \frac{1}{2} + \frac{1}{98} \int_c^{51} \frac{-\log(100 + x) + \log(2 + x)}{x} dx + \frac{1}{98} \int_v^{51} \frac{\log(100 + x) - \log(2 + x)}{x} dx.$$

The proposed mechanism is, by construction, incentive compatible. Hence we can assume truthful reporting.

We need to verify that for any true types in the range, $[2, 100]$ the mechanism is individually rational. For this purpose it is helpful to start of at a good place. Specifically, we choose initial money holdings so that the endow-

ment point lies at the intersection of the two Pareto frontiers that correspond to the case where the seller has the item and the case where the buyer has the item. Set $M = 1$ and choose endowed money holdings m_b and m_s so that utilities are equal if each agent has the mean type. Specifically, set $c = v = \frac{100+2}{2} = 51$ and choose m_b so that

$$51 + \log(1 - m_b) = \log(m_b).$$

Then, $m_b = \frac{e^{51}}{1+e^{51}}$ and $m_s = \frac{1}{1+e^{51}}$. This places us solidly in the interior solution case (the third case in the proof of Theorem 1).

We need to show that for buyer and seller pairs with endowed wealths $m_b = \frac{e^{51}}{1+e^{51}}$ and $m_s = \frac{1}{1+e^{51}}$, and any type profile in $[2, 100]^2$, that both the buyer and the seller are better off under the mechanism than under no trade.

The expected utility of the type c seller under the mechanism is

$$\frac{1}{98} \int_2^{100} \pi(c, v) c + \log\left(\frac{v}{c+v}\right) dv.$$

Why? Since the money allocation of the seller is the same in both states, expected utility is given by the utility of that money allocation plus the expected utility of consuming the item. Both of these components depend upon the type reported by the buyer. Assuming truthful reporting, the seller can compute her expected payoff under the assumption that the buyer's report will be uniformly distributed over the buyer-type range.

Likewise, the expected utility of the type v buyer under the mechanism is

$$\frac{1}{98} \int_2^{100} (1 - \pi(c, v)) v + \log\left(\frac{c}{c+v}\right) dc.$$

The no-trade payoffs for the seller and buyer are $c + \log(\frac{1}{1+e^{51}})$ and $\log(\frac{e^{51}}{1+e^{51}})$, respectively. Hence, for any type $c \in [2, 100]$, the expected net utility gain for the seller under the mechanism is

$$\frac{1}{98} \int_2^{100} \pi(c, v) c + \log\left(\frac{v}{c+v}\right) dv - c - \log\left(\frac{1}{1+e^{51}}\right)$$

and for any type $v \in [2, 100]$ the expected net utility gain for the buyer under the mechanism is

$$\frac{1}{98} \int_2^{100} (1 - \pi(c, v))v + \log\left(\frac{c}{c+v}\right)dc - \log\left(\frac{e^{51}}{1 + e^{51}}\right).$$

The plots in Figure 5 show that both functions are always strictly positive. In particular, the expected net benefit to the seller at $c = 100$ and the buyer at $v = 2$ is $0.7938407 > 0$.




Figure 4: Net utility gain from participating in the mechanism. Left chart shows the gain to the seller for seller types ranging from 2 to 100. Right chart shows the gain to the buyer for buyer types ranging from 2 to 100.

Clearly, the mechanism is most beneficial to low seller types and high buyer types. This makes sense since the gains to trade are greatest when the seller value is the lowest and the buyer value is the highest.

Example 2: CARA preferences

If $U(m) = 1 - e^{-\alpha m}$, then equations (18) and (19) yield $m^B(c, v) = \frac{\ln(c) - \ln(v) + \alpha M}{2\alpha}$ and $m^S(c, v) = \frac{\ln(v) - \ln(c) + \alpha M}{2\alpha}$.²³ Moreover, $\frac{\partial}{\partial c} U(m^S(c, v)) = -e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}}$ and $\frac{\partial}{\partial v} U(m^B(c, v)) = -e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}}$. Substituting these expressions into (20) and (21) yields

$$0 = E_v \left[\frac{\partial \pi(c, v)}{\partial c} c - e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} \right] \quad (11)$$

²³This specification is without loss of generality for $\alpha \geq 0$ and delivers utility values in a range from 0 and 1.

$$0 = E_c \left[-\frac{\partial \pi(c, v)}{\partial v} v - e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} \right] \quad (12)$$

Suppose that c, v are distributed independently and uniformly on $[a, b]$. Define

$$\psi(v) = -E_c \left[-e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} \right] = \frac{1}{b-a} \int_a^b e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} dc = \frac{1}{b-a} e^{-\frac{\alpha M}{2}} \frac{\sqrt{b} - \sqrt{a}}{\sqrt{v}}$$

and

$$\phi(c) = E_v \left[-e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} \right] = -\frac{1}{b-a} \int_a^b e^{-\frac{\alpha M}{2}} \frac{1}{2\sqrt{c}\sqrt{v}} dv = \frac{1}{b-a} e^{-\frac{\alpha M}{2}} \frac{\sqrt{a} - \sqrt{b}}{\sqrt{c}}$$

Following the construction in the proof of Lemma 1, we can set

$$\Delta^s(c, v) = 1, \quad \Delta^b(c, v) = 1.$$

Then, as in Example 1, the functions $b(\cdot)$ and $s(\cdot)$ are given by the ODEs $\psi(v) = vb'(v)$ and $\phi(c) = cs'(c)$, respectively. Thus,

$$b'(v) = \frac{1}{b-a} e^{-\frac{\alpha M}{2}} \frac{\sqrt{b} - \sqrt{a}}{\sqrt[3]{v}}.$$

and

$$s'(c) = \frac{1}{b-a} e^{-\frac{\alpha M}{2}} \frac{\sqrt{a} - \sqrt{b}}{\sqrt[3]{c}}.$$

We use the initial condition $\pi^* = .5$. Then, the seller retains the item with probability .5 when the reports are $\frac{a+b}{2}$. Variations in the seller's report adjusts the probability by the amount

$$\int_c^{\frac{a+b}{2}} s'(x) dx = -\frac{\sqrt{b} - \sqrt{a}}{b-a} e^{-\frac{\alpha M}{2}} \int_c^{\frac{a+b}{2}} \frac{1}{\sqrt[3]{x}} dx$$

and variations in the buyer's report adjusts the probability by the amount

$$\int_v^{\frac{a+b}{2}} b'(x) dx = \frac{\sqrt{b} - \sqrt{a}}{b-a} e^{-\frac{\alpha M}{2}} \int_v^{\frac{a+b}{2}} \frac{1}{\sqrt[3]{x}} dx.$$

So the probability that the seller gets the item if the seller reports c and the buyer reports v is

$$\begin{aligned}
\pi(c, v) &= \frac{1}{2} + \frac{\sqrt{b} - \sqrt{a}}{b - a} e^{-\frac{\alpha M}{2}} \left[\int_v^{\frac{a+b}{2}} \frac{1}{\sqrt[3]{x}} dx - \int_c^{\frac{a+b}{2}} \frac{1}{\sqrt[3]{x}} dx \right] \\
&= \frac{1}{2} + \frac{\sqrt{b} - \sqrt{a}}{b - a} e^{-\frac{\alpha M}{2}} \int_v^c \frac{1}{\sqrt[3]{x}} dx \\
&= \frac{1}{2} + \frac{\sqrt{b} - \sqrt{a}}{b - a} e^{-\frac{\alpha M}{2}} \left[\frac{2}{\sqrt{v}} - \frac{2}{\sqrt{c}} \right].
\end{aligned}$$

The proposed mechanism is, by construction, incentive compatible. Hence we can assume truthful reporting.

To illustrate Theorem 2, we start by setting the range of types equal to $[.3, .7]$ and we choose initial money holdings so that the endowment point lies at the intersection of the two Pareto frontiers that correspond to the case where the seller has the item and the case where the buyer has the item. Set $M = 3$ and choose endowed money holdings m_b and m_s so that utilities are approximately equal if each agent has the mean type (which we will assume is $.5$). Specifically, set $c = v = .5$ and choose m_s so that

$$.5 - e^{-\alpha m_s} \approx -e^{-\alpha(3-m_s)}.$$

Then, $m_s = .535$ and $m_b = 2.465$. Note that it is not necessary to find m_s and m_b that equate the autarky utilities exactly. These values place us solidly in the interior solution case (the third case in the proof of Theorem 1).

We need to show that for buyer and seller pairs with endowed wealths $m_b = 2.465$ and $m_s = .535$, and any type profile in $[.3, .7]^2$, that there exists high enough α such that both the buyer and the seller are better off under the mechanism than under no trade.

It suffices to set $\alpha = 1$. Then, the expected utility of the type c seller under the mechanism is

$$2.5 \int_{.3}^{.7} \pi(c, v) c + 1 - e^{-\frac{3}{2}} \frac{\sqrt{c}}{\sqrt{v}} dv$$

and the expected utility of the type v buyer under the mechanism is

$$2.5 \int_{.3}^{.7} (1 - \pi(c, v))v + 1 - e^{-\frac{3}{2}} \frac{\sqrt{v}}{\sqrt{c}} dc.$$

The no-trade payoffs for the seller and buyer are $c + 1 - e^{-.535}$ and $1 - e^{-2.465}$, respectively. Hence, for any type $c \in [.3, .7]$, the expected net utility gain for the seller under the mechanism is

$$2.5 \int_{.3}^{.7} \pi(c, v)c - e^{-\frac{3}{2}} \frac{\sqrt{c}}{\sqrt{v}} dv - c + e^{-.535}.$$

This function is strictly decreasing in c over the interval $[.3, .7]$ and is positive (equal to 0.19274) at $c = .7$. For any type $v \in [.3, .7]$ the expected net utility gain for the buyer under the mechanism is

$$2.5 \int_{.3}^{.7} (1 - \pi(c, v))v - e^{-\frac{3}{2}} \frac{\sqrt{v}}{\sqrt{c}} dv + e^{-2.465}.$$

This function is strictly increasing in v over the interval $[.3, .7]$ and is positive (equal to 0.021598) at $v = .3$.

Note that if we consider low enough values of α , the preferences become too close to quasilinear and efficient implementation is not possible. For instance, $\alpha = .5$ does not allow efficient trade in the given environment. Then, for seller types close to .7, the mechanism is not individually rational.

Example 3: Shifted Cobb-Douglas preferences

In this example, the seller's type c and the buyer's type v is distributed according to an arbitrary distribution on $[0, 1]$. Moreover, each agent has a shifted Cobb-Douglas utility: $u(x, m, \theta) = (1 + \theta x)m$. The mechanism that generates efficient trade allocates the good and the sum of the money endowments of both agents to the seller with probability $\frac{m_s}{m_s + m_b}$, and it allocates the good and the sum of money endowments to the buyer with the remaining probability, $\frac{m_b}{m_s + m_b}$.

This mechanism ϕ is obviously Bayesian incentive compatible because

the allocation and transfers do not depend on the agents' reports.²⁴ To see that the mechanism is individually rational, notice that the mechanism gives the seller with type c the expected utility of

$$\frac{m_s}{m_s + m_b} (1 + c) (m_s + m_b) = (1 + c) m_s,$$

and that this expected utility is equal to the utility of the seller if no trade takes place. Similarly, the mechanism gives the buyer with type v the expected utility of

$$\frac{m_b}{m_s + m_b} (1 + v) (m_s + m_b) = (1 + v) m_b,$$

and this expected utility is larger than the utility of the buyer if no trade takes place (the latter utility is m_b).

Finally, to see that the mechanism is efficient notice that the Pareto frontier in this example consists of all randomizations among two outcomes: either the seller keeps the good and gets all the financial wealth, or the buyer gets the good and all the financial wealth. This can be immediately seen in Figure 5 (we provide an explicit argument in the appendix).²⁵

It is perhaps not surprising that ex post implementation is possible in this example. Even though agents have private information about their types, private information does not create any uncertainty about the potential for gains from trade. Whomever consumes the good has the highest (constant) marginal utility of wealth and hence efficiency is achieved by giving whomever consumes the good all of the money. Thus, efficiency is achieved by randomizing over the two extreme allocations identified in Figure 5 and the only burden on the mechanism is to define the probability of exchange in a way that ensures individual rationality.

²⁴While mechanism ϕ is also dominant-strategy incentive compatible, our general results will only demonstrate the existence of a Bayesian incentive-compatible mechanism.

²⁵An even simpler example obtains when agents' have standard Cobb-Douglas utilities over the good and the money, $u(x, m; \theta) = A(\theta) x^{\alpha(\theta)} m^{\beta(\theta)}$ for some functions A, α, β . With Cobb-Douglas utilities, the mechanism that allocates the good and all the money to the seller implements efficient trade. Of course, such an example is extreme since, without the indivisible good, agents have no use for money.




Figure 5: Pareto Frontier in the Shifted Cobb-Douglas Example.

7 Concluding Remarks

We focused on providing incentives for agents to truthfully reveal their cost/value information. It is natural to think that preferences are not observable and need to be elicited, while information such as the size of money holdings can be objectively verified. At the same time, in some environments, for instance in the example of Section 5.1, we can not only incentivize agents to reveal their value/cost of the good, we can also provide incentives for them to truthfully announce their money holdings, provided the cost of delivering more money than one has (in the event one is asked to do it) is appropriately high. This is so because—as long as the agent is able to deliver the money—each agent benefits from reporting higher money holdings rather than lower.

Our results on efficient trade open the possibility that other problems might have efficient solutions in non-quasilinear settings. For instance, our results imply the possibility of efficient mechanism for two agents to make

a binary decision, e.g. whether to provide a public good, when each of the agents favors a different decision and each has higher marginal utility of money if his preferred decision is taken.

We establish the generic impossibility of achieving efficient trade in an ex-post incentive compatible way in our setting and we demonstrate that in some special cases efficient trade can be accomplished in strategy-proof way. While we focus on bilateral trade, our analysis can also be used to show that, in the allocation setting, generically no ex-post incentive compatible mechanism is efficient.

Appendix

Derivation of the Pareto Frontier in Section 2

Fix c and v and take any allocation (π, y, z) where π is the probability the seller gets the good, y is the money holding of the seller if she gets the good, and z is the money holding of the seller if he does not get the good. Denoting the total financial wealth by M , the seller's utility is then

$$u^S = \pi(1 + c)y + (1 - \pi)z,$$

and the buyer's utility is

$$u^B = \pi(M - y) + (1 - \pi)(1 + v)(M - z).$$

If $\pi = 0$ or $\pi = 1$, then we are at a deterministic allocation on one of the dotted lines in Figure 1. Consider $\pi \in (0, 1)$. Then, we can weakly increase the utility of both agents by transferring money to the buyer from the seller in the state where the buyer has the good, and transferring the money to the seller from the buyer in the state where the seller has the good. Let us do it so that for every dollar taken from the seller in the state she does not have the good, we allocate her $\frac{1-\pi}{\pi}$ dollars in the state where she has the good, and let us continue transferring the money till either $z = 0$ or $y = M$. Consider the case where we reach the point where $z = 0$ before $y = M$. That

is, we stop at a point where

$$M \geq y + \frac{1-\pi}{\pi} z \quad (13)$$

(the case where we reach $y = M$ before $z = 0$ is symmetric). The utilities after the money transfer are

$$u^{S1} = \pi (1 + c) \left(y + \frac{1-\pi}{\pi} z \right),$$

and

$$\begin{aligned} u^{B1} &= \pi \left(M - y - \frac{1-\pi}{\pi} z \right) + (1-\pi) (1+v) M. \\ &= (M - \pi y - (1-\pi) z) + (1-\pi) v M. \end{aligned}$$

Notice that $u^{S1} \geq u^S$, and thus this transfer provides a Pareto improvement for the seller. Let us now further reallocate the money to the seller in the state she has the good till she has all the money in this state while compensating the buyer through lowering π while keeping the seller's utility constant. Then, the new probability the seller has the good is $\tilde{\pi} = \pi \frac{y + \frac{1-\pi}{\pi} z}{M} = \frac{\pi y + (1-\pi) z}{M}$, and the buyer utility becomes

$$\begin{aligned} u^{B2} &= (1 - \pi') (1+v) M = \left(1 - \frac{\pi y + (1-\pi) z}{M} \right) (1+v) M \\ &= (1+v) (M - \pi y - (1-\pi) z), \end{aligned}$$

which is better than u^B because we are in the case where (13) is satisfied. Thus, both the seller and the buyer prefer a lottery from the postulated Pareto frontier to the initial allocation. To finish the argument notice that no two lotteries from the postulated frontier can be Pareto ranked.

Derivation of (3)

Here we derive the inequalities that play a crucial role in our analysis of the second-order condition. We know from (2) that

$$\frac{\frac{\partial}{\partial m} u(1, m^S, c)}{u(1, m^S, c) - u(0, M - m^B, c)} = \frac{\frac{\partial}{\partial w} u(0, M - m^S, v)}{u(1, m^B, v) - u(0, M - m^S, v)}.$$

Let us show that for $m' \leq m$ (notice that $M - m^B \leq m^S$), the expression $\frac{\frac{\partial}{\partial m} u(1, m, c)}{u(1, m, c) - u(0, m', c)}$ strictly decreases in c . To prove this it is enough to show that

$$\frac{u_c(1, m, c) - u_c(0, m', c)}{u(1, m, c) - u(0, m', c)} > \frac{u_{cm}(1, m, c)}{u_m(1, m, c)}$$

for $m' \leq m$. As in the analysis of SOC, let us rewrite the left-hand side as

$$\begin{aligned} \frac{u_c(1, m, c) - u_c(0, m', c)}{u(1, m, c) - u(0, m', c)} &= \frac{[u_c(0, m, c) - u_c(0, m', c)] + [u_c(1, m, c) - u_c(0, m, c)]}{[u(0, m, c) - u(0, m', c)] + [u(1, m, c) - u(0, m, c)]} \\ &= \frac{[u_c(1, m, c) - u_c(0, m, c)] + \int_{m'}^m u_{cm}(0, \tilde{m}, c)}{[u(1, m, c) - u(0, m, c)] + \int_{m'}^m u_m(0, \tilde{m}, c)}. \end{aligned}$$

Now, (1) gives

$$\frac{u_c(1, m, c) - u_c(0, m, c)}{u(1, m, c) - u(0, m, c)} > \frac{\frac{\partial}{\partial m} u_c(0, m, c)}{\frac{\partial}{\partial m} u(0, m, c)},$$

and the constancy of $\frac{u_{cm}(0, \tilde{m}, c)}{u_m(0, \tilde{m}, c)}$ in \tilde{m} gives

$$\frac{\frac{\partial}{\partial m} u_c(0, \tilde{m}, c)}{\frac{\partial}{\partial m} u(0, \tilde{m}, c)} = \frac{\frac{\partial}{\partial m} u_c(0, m, c)}{\frac{\partial}{\partial m} u(0, m, c)}.$$

Together this demonstrates that $\frac{\frac{\partial}{\partial m} u(1, m, c)}{u(1, m, c) - u(0, m', c)}$ strictly decreases in c .

Going back to our initial equality, as we increase c while keeping money levels constant the left-hand side decreases, while the right hand side stays constant. Looking at the graph shows that to balance this out, the critical point shifts towards higher utility of the buyer. And, thus the seller's money level m^S decreases in this critical point. The analysis of other critical money levels is similar.

Proof of Lemma 1

We will develop a constructive procedure to find proper randomization. As a preparation, consider the PDE

$$S_1(c, v) \frac{\partial}{\partial c} \pi(c, v) + S_2(c, v) \pi(c, v) = 0, \quad (14)$$

$$B_1(c, v) \frac{\partial}{\partial v} \pi(c, v) + B_2(c, v) \pi(c, v) = 0, \quad (15)$$

Considered separately, these equations are standard ODEs. They have solutions, and on a bounded domain we can assume that the solutions are positive. We can thus fix a solution $\Delta^b > 0$ to the first equation and a solution $\Delta^s > 0$ to the second. Consider functions $b(\cdot)$ and $s(\cdot)$, and set

$$\pi(c, v) = b(v) \Delta^b(c, v) + s(c) \Delta^s(c, v).$$

Consider the second PDE equation from the lemma, and notice that the first summand above is zero for each v , and thus it is zero in expectation. Thus, the second equation reduces to

$$\begin{aligned} \psi(v) &= E_c \left[B_1(c, v) \frac{\partial}{\partial v} \left[b(v) \Delta^b(c, v) \right] + B_2(c, v) b(v) \Delta^b(c, v) \right] = \\ &= E_c \left[B_1(c, v) \Delta^b(c, v) \right] b'(v) + E_c \left[B_1(c, v) \frac{\partial}{\partial v} \Delta^b(c, v) + B_2(c, v) \Delta^b(c, v) \right] b(v) \end{aligned}$$

Since, $B_1 \Delta^b > 0$ this equation has solutions. Let b be one such solution satisfying the initial condition $b(v^*) \Delta^b(c^*, v^*) = \frac{1}{2} \pi^*$. Similarly, we can find function s for which the first PDE equation from the lemma is satisfied, and such that $s(c^*) \Delta^s(c^*, v^*) = \frac{1}{2} \pi^*$. Thus, for these two functions b and s , the function π defined above satisfies the system of PDE from the lemma, as well as the initial condition. The flatness claim of the lemma now follows because Δ^s and Δ^b are nearly flat (since the coefficient in front of the derivative is separated from zero, and the other parts of the equation are close to zero), and because b and s are nearly flat (for the same reason). QED

Derivation of the efficient contract for the case of separable utilities

In the separable case, the equations given in (2) that define the money levels $m^S(c, v)$ and $m^B(c, v)$ associated with the points S and B in Figures 2-4, respectively, reduce to

$$\frac{\frac{\partial}{\partial m} U(m^S)}{\frac{\partial}{\partial m} U(M - m^S)} = \frac{\frac{\partial}{\partial m} U(M - m^B)}{\frac{\partial}{\partial m} U(m^B)} \quad (16)$$

and

$$c + U(m^S) = U(M - m^B) + \frac{\frac{\partial}{\partial m} U(m^S)}{\frac{\partial}{\partial m} U(M - m^S)} * [v + U(m^B) - U(M - m^S)]. \quad (17)$$

Equation (16) implies

$$m^S = M - m^B. \quad (18)$$

To see this, suppose $m^S > M - m^B$. Then, $m^B > M - m^S$ and since U is strictly increasing in m , the LHS of (16) would be strictly greater than the RHS, contradicting (18). The reverse contradiction occurs if we assume $m^S < M - m^B$. In other words, in the optimal contract for the separable case, each player has equal money in each state.

Substituting (18) into (17) yields

$$c = \frac{\frac{\partial}{\partial m} U(m^S)}{\frac{\partial}{\partial m} U(M - m^S)} v. \quad (19)$$

Given any pair (c, v) , equation (19) uniquely defines $m^S(c, v)$ and $m^B(c, v)$ is then given by (18).

The relation $m^S(c, v) = M - m^B(c, v)$ implies that $\frac{\partial}{\partial \hat{c}} u(1, m^S(\hat{c}, v); c) = \frac{\partial}{\partial \hat{c}} u(0, M - m^B(\hat{c}, v); c) = \frac{\partial}{\partial \hat{c}} U(m^S(\hat{c}, v))$ and that $\frac{\partial}{\partial \hat{v}} u(0, M - m^S(c, \hat{v}); v) = \frac{\partial}{\partial \hat{v}} u(1, m^B(c, \hat{v}); v) = \frac{\partial}{\partial \hat{v}} U(m^B(c, \hat{v}))$. Thus, the first-order equations (5) and (6) that define $\pi(c, v)$ the incentive compatible contract become

$$0 = E_v \left\{ \frac{\partial \pi(c, v)}{\partial c} c + \frac{\partial U(m^S(c, v))}{\partial c} \right\}. \quad (20)$$

and

$$0 = E_c\left\{-\frac{\partial\pi(c, v)}{\partial v}v + \frac{\partial U(m^B(c, v))}{\partial v}\right\}, \quad (21)$$

respectively.

Equation (1) is satisfied in the separable case since $\frac{\partial \log(\theta)}{\partial \theta} = \frac{1}{\theta} > 0$ and $\frac{\partial \log(U'(m))}{\partial \theta} = 0$. We established in Section 3.1.2 that equation (1) is sufficient for the second-order conditions of the incentive-compatibility analysis to hold. Hence, the second-order conditions are satisfied in the separable case.

References

- Baisa, Brian (2013). "Auction Design without Quasilinear Preferences." working paper.
- Chatterjee, K., and W. Samuelson(1983). "Bargaining under Asymmetric Information." *Operations Research* 31, 835–851.
- Coase R. (1960). "The Problem of Social Cost," *Journal of Law and Economics* 3, 1–44.
- Cook, Philip J. and Daniel A. Graham (1977). "The Demand for Insurance and Protection: The Case of Irreplaceable Commodities." *Quarterly Journal of Economics* 91, 143–156.
- Cramton, Peter, Robert Gibbons, and Paul Klemperer (1987). "Dissolving a Partnership Efficiently." *Econometrica* 55(3), 615–632.
- Crawford, Vincent P. (2016) "Efficient Mechanisms for Level-k Bilateral Trading."
- Cremer, Jacques and Richard P. McLean (1988). "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56(6), 1247–57.
- Cripps, Martin. W. and Swinkels, Jeroen. M. (2006). "Efficiency of Large Double Auctions." *Econometrica* 74(1), 47–92.
- Garratt, Rod (1999). "On Bargaining for an Indivisible Good." *Games and Economic Behavior* 26, 186–192.
- Gresik, Thomas A., and Mark A. Satterthwaite, 1983. "The Number of

- Traders Required to Make a Market Competitive: The Beginnings of a Theory," Discussion Papers 551, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- Holt, Charles. 1980. "Competitive Bidding for Contracts under Alternative Auction Procedures," *J. Polit. Econ.* 88(3), 433–45.
- Kos, Nenad, and Mihai Manea, (2009) Efficient Trade Mechanism with Discrete Values." Working Paper.
- Krishna, Vijay and Motty Perry (1997). "Efficient Mechanism Design." Unpublished paper, Penn State University.
- Loertscher, Simon, Leslie M. Marx, and Tom Wilkening, (2013). "A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers." Working Paper.
- Makowski, Louis, and Joseph M. Ostroy, 1989. "Efficient and Individually Rational Bayesian Mechanisms Only Exist on Perfectly Competitive Environments." *UCLA Economics Working Papers 566*, UCLA Department of Economics.
- Mardare, Sorin (2007) "On Systems of First Order Linear Partial Differential Equations with L^p Coefficients." *Advances in Differential Equations* 12, 301-360.
- Maskin, Eric and John Riley (1984). "Optimal Auctions with Risk Averse Buyers." *Econometrica* 52(6), 1473-1518.
- Matthews, Steven. 1983. "Selling to Risk Averse Buyers with Unobservable Tastes," *J. Econ. Theory* 30, 370–400.
- Matsuo, T. (1989): "On Incentive Compatibility, Individual Rationality, and Ex Post Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory*, 49, 189-194.
- McAfee, R. Preston and Philip J. Reny (1992). "Correlated Information and Mechanism Design." *Econometrica* 60(2) 395-421.
- McAfee, R. Preston (1991). "Efficient Allocation with Continuous Quantities." *Journal of Economic Theory* 53 (1), 51-74.
- Milgrom, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press.
- Myerson, Roger B. and Mark A. Satterthwaite (1983). "Efficient Mechanisms

- for Bilateral Trading." *Journal of Economic Theory* 29 (2), 265-281.
- Reny, Philip J. and Motty Perry (2006). "Toward a Strategic Foundation for Rational Expectations Equilibrium." *Econometrica* 74: 1231–1269.
- Riley, John (2012). *Essential Microeconomics*. Cambridge University Press.
- Rustichini, Aldo, Mark A. Satterthwaite and Steven R. Williams (1994) "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica* 62(1), 1041—1063.
- Thomas, Tracey Y. (1934). "Systems of Total Differential Equations Defined Over Simply Connected Domains." *Annals of Mathematics* 35 (4), 730-734.
- Williams, Steven R. (1999). "A Characterization of Efficient Bayesian Incentive Compatible Mechanisms." *Economic Theory* 14, 155-180.
- Wilson, Robert (1985). "Incentive Efficiency of Double Auctions." *Econometrica* 53, 1101-1115.
- Wolitzky, Alexander (2015). "Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade."