

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

The role of commitment in bilateral trade

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1723676> since 2020-01-16T19:08:39Z

Published version:

DOI:10.1016/j.jet.2014.09.017

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

The Role of Commitment in Bilateral Trade*

Dino Gerardi[†], Johannes Hörner[‡], Lucas Maestri[§]

June 27, 2014

Abstract

This paper characterizes the set of equilibrium payoffs in bargaining with interdependent values when the informed party makes all offers, as discounting vanishes. The seller of a good is informed of its quality, which affects both his cost and the buyer's valuation, but the buyer is not. To characterize this payoff set, we derive an upper bound, using mechanism design with limited commitment. We then prove that this upper bound is tight, by showing that all its extreme points are equilibrium payoffs. Our results shed light on the role of different forms of commitment in bargaining. In particular, they imply that the buyer's inability to commit before observing the terms of trade is what precludes efficiency.

Keywords: bargaining, mechanism design, market for lemons.

JEL codes: C70, C78, D82

1 Introduction

With few exceptions, non-cooperative theories of bargaining concern themselves with the extreme cases of full commitment or no commitment whatsoever. Mechanism design neglects

*We thank an Associate Editor and two referees, as well as Joel Watson for useful comments, Sofia Moroni for careful proofreading, and various seminar audiences.

[†]Collegio Carlo Alberto, Torino, Italy. dino.gerardi@carloalberto.org.

[‡]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA. johannes.horner@yale.edu.

[§]FGV/EPGE – Escola Brasileira de Economia e Finanças, Praia de Botafogo, 190/1126C, Rio de Janeiro, RJ, Brazil. lucas.maestri@fgv.br.

the problem of sequential rationality, whereas infinite-horizon bargaining following Rubinstein (1982) reveals in the asymptotic analysis as frictions, interpreted as commitment, disappear.

Our goal is to better understand how commitment matters in markets with adverse selection. To do so, we consider a standard model of trade, with one buyer and one seller. The setting is that of the lemon problem, as introduced by Akerlof (1970), the simplest framework for trade under interdependent values. The seller knows both the value and cost of the unit, while the buyer does not.¹ There is common knowledge of gains from trade.

The full commitment problem has been thoroughly investigated by Samuelson (1984) and Myerson (1985). As already pointed out by Myerson (1981), optimal mechanisms with interdependent values exhibit surprising properties. In particular, the optimal mechanism need not satisfy “posterior” individual rationality. The buyer may lose from participating in the mechanism given the information that this mechanism reveals: if the buyer were to reconsider his willingness-to-trade in light of the offer that he is meant to accept, he may prefer to pass.

Giving such veto power to the buyer is the second step in our analysis. Note that this is not equivalent to *ex post* individual rationality, a stronger requirement that posits that the buyer gains given the actual state of nature. The difference matters here, since values are interdependent (see Gresik 1991b, Forges 1994 and Matthews and Postlewaite 1989). This property, which we refer to as *veto-incentive compatibility*, following Forges (1999), imposes restrictions on the mapping from reported types to the distribution over offers that the mechanism specifies. Veto-incentive compatibility, then, is a requirement on the graph of this map: conditional on any given offer, the posterior belief of the buyer should be such that he is willing to accept this offer. Restricting attention to deterministic offers would entail a loss of generality. This does not mean that the buyer accepts a random price; rather, the price he accepts is chosen randomly.

Veto-incentive compatibility is not only a restriction that is realistic, given current legal and commercial practices, it is also *implied* by standard bargaining protocols. We prove that it is

¹By a simple change of variable, all our results apply to the case in which it is the buyer who is informed and who makes offers, and the seller is uninformed.

automatically satisfied whenever at most one of the players makes an offer *per* round, whether or not this order is stochastic, history-dependent, etc., and whether the horizon is finite or not.

We characterize veto-incentive compatible mechanisms. We prove that whether a given allocation can be implemented in a veto-incentive compatible way is a property of (the map from reports to) the probability of trade and expected price alone. The problem reduces to standard optimal control. The interesting feature is the restriction implied by veto-incentive compatibility: the necessary and sufficient condition is that the buyer's *ex ante* payoff, conditional on trading with all types above a given threshold be nonnegative, for all possible values of this threshold.

Finally, we consider the case of no commitment, as captured by infinite-horizon bargaining. We focus on the case in which the seller makes all offers. The resulting mechanism must satisfy veto-incentive compatibility, since the buyer can reject any given offer. The temporal monopoly of the seller provides him with a lower bound on his payoff. Namely, he can secure a price equal to the buyer's lowest possible value. We prove that, along with veto-incentive compatibility, this is the only further constraint imposed by bargaining: every payoff vector that can be achieved by a veto-incentive compatible allocation and that gives the seller this security payoff is an equilibrium payoff vector if the two players are patient enough. This might sound like a folk theorem, but this only holds in terms of payoffs: there are allocations that are veto-incentive compatible, and give the seller his security payoff, and yet cannot be implemented in the bargaining game.

Our results have striking implications. First, under the sufficient conditions we provide, bargaining achieves constrained efficiency. In those cases, commitment has no benefit whatsoever. Second, if bargaining fails to attain efficiency, then trading institutions are only useful if they manage to weaken the veto-incentive compatibility constraint, as is the case, for instance, when the uninformed party is asked to commit to a screening contract. Nothing in between helps.

Our results are only about *ex ante* payoffs: Sequential rationality imposes further constraints. For instance, both the seller's and the buyer's payoffs must be individually rational, not only *ex ante*, but from any history onward. Some veto-incentive compatible allocation that give the seller

his security payoff cannot be implemented in the bargaining game. This “folk” theorem does not extend to interim payoffs, and might not include the equilibrium payoff of the game in which the buyer makes all offers (Deneckere and Liang, 2006). Surprisingly, these other constraints do not affect the set of *ex ante* payoffs. Together with Deneckere and Liang, this paper clarifies the role of the proposer’s identity. For instance, the most efficient equilibrium outcome when the seller makes all offers is more efficient than the equilibrium outcome when the buyer makes offers. Even the most inefficient equilibrium might be better.

Among related papers, Ausubel and Deneckere (1989) analyze the link between mechanism design and bargaining in the special case of private values (with one-sided incomplete information). They show that, when the uninformed party makes all the offers, a folk theorem holds. On the other hand, if the informed party makes the offer, a unique equilibrium outcome gets singled out as the frequency of offers increases. Our paper establishes that, as one would suspect, lack of commitment imposes more constraints with interdependent values than with private values. Interestingly though, the set of equilibrium payoffs that can be achieved remains fairly easy to characterize, as the feasible set of a static optimization program. The paper by Deneckere and Liang that was already mentioned provides a careful analysis of the bargaining game in which the (uninformed) buyer makes all offers, and they prove that the equilibrium outcome is unique.

Section 2 defines the set-up. The main results are in Section 3, with a sketch of proof provided in Section 4. (The full proof is available in an online appendix.) Section 5 offers extensions.

2 The Set-Up

2.1 The Trading Problem

Consider a trading problem in which player 1, the *seller*, owns an indivisible object that player 2, the *buyer*, wants to purchase. The two players are risk-neutral, with quasi-linear utility. The players’ valuations are determined by the realization of a random variable that is uniformly

distributed over the unit interval, $t \sim \mathcal{U}[0, 1]$. That is, given t , the seller's cost and the buyer's value for the object are given by $c(t)$ and $v(t)$, respectively. The functions $c : [0, 1] \rightarrow \mathbb{R}_+$ and $v : [0, 1] \rightarrow \mathbb{R}_+$ are assumed to be non-decreasing, right-continuous and piecewise C^1 . Because v need not be constant, this environment displays interdependent values, of which private values is a special case. The assumption that t is uniformly distributed is made with no loss of generality, given the restrictions imposed on v and c .²

Information is asymmetric. The seller is informed of the realization of the random variable, and so knows both his cost and the buyer's value. We refer to this realization as the seller's *type* $t \in T := [0, 1]$. The buyer, on the other hand, does not observe this realization. However, he knows the distribution of the random variable, and the functions v and c are common knowledge.

In particular, it is common knowledge that there are gains from trade. That is, we assume that $v(t) > c(t)$.³ This neither precludes nor implies that the first-best allocation is attainable if individual rationality is imposed. Such a first-best mechanism is individually rational if and only if the buyer's expected value exceeds the seller's highest cost (see Lemma 1 of Deneckere and Liang, 2006). While our results can be adapted to this case, the trading problem becomes then uninteresting, and we will rule it out.

Our purpose is to characterize the equilibrium payoffs in the bargaining game in which the seller makes all offers. To do so, we must understand what allocations can be achieved under limited commitment. First, we shall consider the case in which the buyer cannot be forced to trade if the actual offer that is being made leads to a negative expected payoff. Following Forges (1999), we refer to this assumption as *veto-incentive compatibility*. Given the mechanism, and for any outstanding offer, the buyer updates his expected value for the object. Veto-incentive

²In an online appendix (Appendix C), we discuss what happens when c and v are not both monotonic. Sufficiency from Proposition 1 survives, and necessary and sufficient conditions are provided under which the *ex ante* efficient payoff in the full commitment is an equilibrium payoff of the bargaining game as frictions vanish.

³The main results –Theorem 1, Proposition 1, Proposition 2, Theorems 2 and 3– also hold under the weaker condition $v \geq c$, as long as there are finitely many types. Our proofs involve a series of approximation of the cost and value functions by step functions, and we have not verified that the limit results extend to the case of a weak inequality.

compatibility requires this conditional expectation to exceed the offer, whenever the mechanism specifies trade in this event. This captures the notion that, in most trading environments, buyers can always reject an offer for which they anticipate a loss. In the words of Gresik (1991a), “in most markets each trader has the ability to refuse to trade when the “best” negotiated terms give him negative utility.” For instance, a seller who puts up an object for sale in an auction house commits to the eventual outcome, given the auction mechanism, but potential buyers can drop out at any stage of the auction process. Note that, with interdependent values, this does not ensure that the buyer will not experience regret, that is, that his realized value will exceed the price that he paid. In many markets, there is not much a buyer can do to renege on a purchase for which his experienced utility falls short of the price that he paid. In this sense, the trade need not be *ex post* individually rational. (The two notions coincide in the case of private values.) At the time of purchase, however, the potential buyer cannot be forced to accept an outstanding offer, if he anticipates a loss, simply because he chose to participate in the trading process.

The set of payoffs that can be achieved under this mechanism (as well as under the standard “full-commitment” mechanism) will then be compared to the set of payoffs in the infinite-horizon bargaining game with discounting, in which the seller makes all the offers.

2.2 Mechanisms

Direct mechanisms, that require the seller to report his type, provide a way for setting the terms of trade. To be more formal, a *direct mechanism* is a probability transition μ from T to $\{0, 1\} \times \mathbb{R}_+$.^{4,5} A direct mechanism, then, specifies whether trade occurs (the outcome “1” is interpreted as trade, while the outcome “0” means no trade), and at what price, according to some joint distribution, given the announcement of the seller. We let $x(t)$ denote the probability

⁴That is, for each $t \in T$, $\mu(t)$ is a probability distribution on $\{0, 1\} \times \mathbb{R}_+$, and the probability $\mu(\cdot)[A]$ assigned to any Borel set $A \subset \{0, 1\} \times \mathbb{R}_+$ is a measurable function of $t \in T$. That attention can be restricted to distributions over the probability of trade and payment is a consequence of the revelation principle.

⁵It is not hard to see that the restriction to offers in \mathbb{R}_+ rather than \mathbb{R} is without loss of generality for those allocations, and hence payoffs, that we seek to characterize.

of trade, given the announcement t . That is,

$$x(t) := \mu(t)[1, \mathbb{R}_+]. \quad (1)$$

Without loss of generality, we assume that no payment is made if no trade occurs, that is, we assume that $\mu(t)[0, \{0\}] = 1 - x(t)$. If $x(t) > 0$, we let $p(t)$ denote the expected price, given the announcement t , *i.e.*

$$p(t) := \int_{\mathbb{R}_+} p \mu(t)[1, dp] / x(t), \quad (2)$$

and set $p(t) := 0$ otherwise. Given $x : T \rightarrow [0, 1]$ and $p : T \rightarrow \mathbb{R}_+$, the *allocation* (x, p) is *implementable* if there exists a mechanism μ (which *implements* (x, p)) such that x and p coincide everywhere with the functions that are defined by (1) and (2).

It follows from the revelation principle that attention can be restricted to direct mechanisms in which the seller announces his type truthfully. Furthermore, under commitment, attention can be restricted to mechanisms in which prices are deterministic, *i.e.* $p(t)$ is the only price assigned positive probability by $\mu(t)[1, \cdot]$, for all t .

Given some direct mechanism μ , the payoff to the seller of type t that reports s is given by

$$\pi^S(s | t) := x(s)[p(s) - c(t)].$$

The mechanism μ is *incentive compatible* if, for all $s, t \in T$, $\pi^S(t) := \pi^S(t | t) \geq \pi^S(s | t)$. We shall also be interested in the *ex ante* payoff of the seller before his type is determined, that is, given some incentive compatible mechanism μ ,

$$\pi^S = \mathbb{E}_t[\pi^S(t)] = \int_T \pi^S(t) dt = \int_T x(t)[p(t) - c(t)] dt. \quad (3)$$

Fix some incentive compatible mechanism μ . Suppose that the buyer is offered to trade at some price p in the support of $\mu(t)[1, \cdot]$ for some $t \in T$. What is his expected payoff, conditional on this outcome $(1, p)$? Given the mechanism μ , fix a version of the conditional distribution $\nu : (\{0, 1\} \times \mathbb{R}_+) \times \mathcal{B} \rightarrow [0, 1]$, where \mathcal{B} is the Borel field on T . Given $\mathcal{T} \in \mathcal{B}$, we write $\nu(\mathcal{T} | p)$

for $\nu((1, p), \mathcal{T})$, the conditional probability assigned to the seller's type being in the set \mathcal{T} , given the event $(1, p)$ (with an abuse of notation, we also write $\nu(t | p)$ for $\nu(\{t\} | p)$). The buyer's expected payoff, given p , is then

$$\pi^B(p) := \int_T v(t) d\nu(t | p) - p.$$

The *ex ante* payoff of the buyer is given by

$$\pi^B := \int_T x(t)[v(t) - p(t)] dt. \quad (4)$$

An incentive compatible mechanism μ is *individually rational* if $\pi^S(t) \geq 0$ for all $t \in T$, and $\pi^B \geq 0$. It is *veto-incentive compatible* if $\pi^B(p) \geq 0$ for all prices in the support of μ . Because the buyer must break even given his conditional expectation, one cannot restrict attention to mechanisms specifying deterministic prices, when considering veto-incentive compatible mechanisms.

To summarize, we shall be interested in determining the allocations (x, p) that can be implemented by incentive compatible, individually rational and veto-incentive compatible mechanisms, and in the set of *ex ante* payoffs $\pi = (\pi^B, \pi^S)$ spanned by such allocations.⁶ For short, we refer to this problem as the veto-incentive compatible program, and these allocations as the veto-incentive compatible allocations, to be compared with the *full commitment allocations*, in which the requirement of veto-incentive compatibility is dropped. The problem of determining the latter set is known (see, in particular, Samuelson 1984, and Myerson 1985), and is referred to in the sequel as the full commitment program.

Of particular interest is the (constrained) *efficient* allocation for each program, that is, any allocation (x, p) that maximizes the overall gains from trade $\int_T x(t)[v(t) - c(t)] dt$, or equivalently, that maximizes the sum of *ex ante* payoffs $\pi^S + \pi^B$.

⁶A set of allocations $\{(x, p)\}$ *spans* the payoff set $A \subset \mathbb{R}^2$ if the image of that set, by the mappings defined by (3) and (4), is equal to A .

2.3 The Bargaining Game

In Section 3.3, we shall finally consider the infinite-horizon bargaining game. Trivially, this further reduces the set of implementable allocations. Deneckere and Liang (2006) have provided a comprehensive analysis of the game in which the uninformed party, the buyer, makes all the offers. Doing so allows to abstract from signaling issues: given any history, there is only one action available to the informed party which does not terminate the game. Therefore, the analysis becomes tractable, although far from trivial, and the equilibrium outcome turns out to be unique. We shall consider the opposite case, in which the seller makes all the offers, and show that, in this case as well, it is possible to provide a simple characterization of the equilibrium payoffs as bargaining frictions vanish. Furthermore, the best equilibrium improves upon the equilibrium in the game in which the buyer makes the offers (in terms of efficiency).

Let us define the game formally. Time is discrete, and indexed by $n = 1, \dots, \infty$. At each time or period n , the seller asks a price for the unit. After observing the price, the buyer either accepts or rejects it. If the price is accepted, the game ends. If the offer is rejected, a period elapses and the seller asks for a price again. We shall allow for a public randomization device in the initial period (for concreteness, think of a draw from the uniform distribution on the unit interval), before the seller sets the first price. This allows us to focus on the extreme points of the equilibrium payoff set, and we shall not refer to this randomization device in the sequel.

The seller's asking price can take any real value. An outcome of the game is a triple (t, n, p_n) , with the interpretation that the realized type is t , and that the buyer accepts the seller's price p_n in period n (which implies that all previous prices were rejected). The case $n = \infty$ corresponds to the outcome in which the buyer rejects all the prices (as a convention, set p_∞ equal to 0). Buyer and seller discount future payoffs at the common discount factor $\delta \in (0, 1)$. The seller's von Neumann-Morgenstern utility function over outcomes is his net surplus $\delta^{n-1}(p_n - c(t))$ when $n < \infty$, and zero otherwise. This suggests the interpretation of the cost as an actual production cost incurred at the time of the transaction, but an alternative and equivalent formulation is that

the seller derives a flow utility of $(1 - \delta)c(t)$ in every period in which he holds on to the unit.

The buyer's utility is $\delta^{n-1}(v(t) - p_n)$ when the outcome is (t, n, p_n) , $n < \infty$, and 0 if $n = \infty$.⁷ The players' expected utilities over lotteries of outcomes, or *payoffs*, are defined as usual.

A history (of prices) $h^{n-1} \in H^{n-1}$ in case trade has not occurred by time n is a sequence (p_1, \dots, p_{n-1}) of asking prices that the seller set and the buyer rejected (set $H_0 := \emptyset$). A behavior strategy σ^S for the seller is a sequence $\{\sigma_n^S\}$, where σ_n^S is a probability transition from $T \times H^{n-1}$ into \mathbb{R} , mapping the seller's type, the history h^{n-1} into a (possibly random) asking price. A behavior strategy σ^B for the buyer is a sequence $\{\sigma_n^B\}$, where σ_n^B is a probability transition from $H^{n-1} \times \mathbb{R}$ into $\{0, 1\}$, mapping the history h^{n-1} and the outstanding price into a probability of acceptance (as before, "1" denotes acceptance, and "0" rejection). We use the perfect Bayesian equilibrium (PBE) concept as defined in Fudenberg and Tirole (1991, Definition 8.2).⁸ Given some (perfect Bayesian) equilibrium, we follow standard terminology in calling a seller's offer *serious* if it is accepted by the buyer with positive probability. An offer is *losing* if it is not serious. Clearly, the specification of losing offers in an equilibrium is, to a large extent, arbitrary.

Given some equilibrium $\sigma = (\sigma^B, \sigma^S)$, we denote by $\pi^S(\sigma)$ and $\pi^B(\sigma)$ the *ex ante* payoff of the seller and the buyer, respectively. Note that this involves taking expectations with respect to the seller's type. Given δ , the payoff vector $\pi = (\pi^B, \pi^S)$ can be *achieved* in the bargaining game if there exists an equilibrium σ of the bargaining game such that $\pi = (\pi^B(\sigma), \pi^S(\sigma))$.

Let $E(\delta)$ denote the set of equilibria in the bargaining game with discount factor δ , and $\Pi(\delta) \subset \mathbb{R}^2$ the set of payoff vectors given δ . Further, define $\underline{\Pi} := \liminf_{\delta \rightarrow 1} \Pi(\delta)$ and $\overline{\Pi} := \limsup_{\delta \rightarrow 1} \Pi(\delta)$ as the inner and outer limits of the equilibrium payoff set as frictions vanish. We shall show that those two sets are equal, and provide a simple characterization of this set.

⁷Discounting plays no role for optimality. Results carry over to a sequence of short-run buyers, as long as the buyer's payoff is interpreted as the discounted sum of these short-run buyers' payoffs.

⁸Fudenberg and Tirole define perfect Bayesian equilibria for finite games of incomplete information only. The suitable generalization of their definition to infinite games is straightforward and omitted.

3 Main Results

3.1 Preliminaries: The Full Commitment Program

We start by recalling the characterizations obtained by Samuelson (1984) and Myerson (1985) for the set of *ex ante* payoffs that can be achieved through mechanisms that satisfy incentive compatibility and individual rationality.

For later purposes, it is useful to define the following. Given a mechanism μ , define the expected payment $\bar{p}(t)$ received by type $t \in T$ as

$$\bar{p}(t) := x(t)p(t).$$

Note that specifying the function $\bar{p} : T \rightarrow [0, 1]$ is equivalent to specifying the function p , given our convention that $p(t) = 0$ whenever $x(t) = 0$. Incentive compatibility is the requirement that

$$\pi^S(t) = \bar{p}(t) - x(t)c(t) \geq \bar{p}(s) - x(s)c(t),$$

for all $s, t \in T$. This implies, in particular, that

$$\pi^S(t) \geq \lim_{s \downarrow t} \pi^S(s | t),$$

for all $t \in T$. We refer to these constraints as the set of *local incentive compatibility constraints*.

Suppose that the local incentive compatibility constraints are binding for all $t \in T$.⁹ It is then standard to show that π^S has bounded variation and equal to, for all t ,¹⁰

$$\pi^S(t) = \pi^S(1) + \int_t^1 x(s)dc(s).$$

In this case, all expected payments are uniquely determined by the probabilities of trade (and

⁹Because the cost function need not be continuous, there are allocations that are implementable in the full commitment program for which some local incentive compatibility constraints are not binding.

¹⁰Here and in what follows, $\int_T x(s)dc(s) := \int_{(0,1)} x(t)c'(t)dt + \sum_{t \in D^c} x(t)(c(t) - \lim_{s \uparrow t} c(s))$, where c' is the derivative of c on each interval, D^c is the set of discontinuities of c , and x is assumed to be right-continuous (since c and v are, this is without loss of generality). Later references to derivatives have to be understood similarly.

the price $\bar{p}(1)$) through

$$\bar{p}(t) = \bar{p}(1) - x(1)c(1) + x(t)c(t) + \int_t^1 x(s)dc(s).$$

Let us also define the buyer's payoff $B(t)$ accruing from all seller's types above t , given some allocation (x, p) , as

$$B(t) := \int_t^1 (x(s)v(s) - \bar{p}(s))ds. \quad (5)$$

Note that $B(0) = \pi^B$. Further, if all local incentive compatibility constraints are binding, we can express $B(t)$ as a function of x (and $\bar{p}(1)$) only. Explicitly,

$$\begin{aligned} B(t) &= \int_t^1 [x(s)v(s) - (\bar{p}(1) - x(1)c(1) + x(s)c(s) + \int_s^1 x(u)dc(u))]ds \\ &= \int_t^1 [x(s)(v(s) - c(s)) - \int_s^1 x(u)dc(u)]ds - (1-t)(\bar{p}(1) - x(1)c(1)). \end{aligned}$$

Trivially, given the revelation principle, the set of implementable allocations in the full commitment program is characterized by incentive compatibility and individual rationality. A sharper characterization can be obtained for the set of achievable payoff vectors.

Theorem 1 (*Samuelson 1984, Myerson 1985*)¹¹ *Assume $c(1) \geq \int_T v(t)dt$. Under full commitment:*

1. *The payoff set can be obtained, without loss of generality, by assuming that all local incentive compatibility constraints bind, and that the highest seller type's payoff is zero: $\pi^S(1) = 0$;*
2. *The payoff set is spanned by the set of non-increasing functions $x : T \rightarrow [0, 1]$ subject to*

$$\int_0^1 [x(s)(v(s) - c(s)) - \int_s^1 x(u)dc(u)]ds \geq 0,$$

given expected payments, for all $t \in T$,

¹¹More precisely, Samuelson (1984, unnumbered lemma) shows that the Pareto-efficient allocations are achieved by two- or three-step functions. Parts 1-2 follows from Myerson's analysis (although all conclusions are rather straightforward given Samuelson's result).

$$\bar{p}(t) = x(t)c(t) + \int_t^1 x(s)dc(s).$$

3. *The payoff set is a convex polygon whose extreme points are achieved by functions $x : T \rightarrow [0, 1]$ that are step functions with either two or three steps; the origin is an extreme point, and for all other extreme points, it can be assumed that $x(0) = 1$.*

The constraint in the second part of the theorem is simply the requirement that $B(0) \geq 0$, given the definition of \bar{p} . The requirement that x be non-increasing ensures incentive compatibility, given the definition of \bar{p} . Theorem 1.2 states that any non-increasing function $x \in [0, 1]$ satisfying $B(0) \geq 0$ (a constraint that only involves x) is part of an allocation that is implementable in the full commitment program, along with the expected payments defined in the theorem, and that these allocations are a sufficient class to generate all the payoffs that can be achieved in this program. As mentioned, one mechanism implementing any such allocation is a mechanism with deterministic prices. There are other mechanisms implementing this allocation, and other allocations that are implementable, but they do not lead to any extra payoff vectors.

In light of this characterization, the payoff set of the full commitment program can be obtained by considering a family of continuous linear programs, in which one maximizes $\lambda \cdot \pi$ over functions x satisfying the constraints given in Theorem 1.2, where $\lambda \in \mathbb{R}^2$ are the (possibly negative) weights on the buyer and seller's payoffs. The maxima of these programs determine the extreme points of the payoff set, and it is then a standard result that such extreme points are themselves achieved by extreme points of the admissible set, *i.e.*, by step functions.

The (constrained) efficient allocation takes a simple form, given that it solves a maximization problem in which the objective and the constraint are linear. As Samuelson and Myerson show, the *ex ante* efficient mechanism is as follows: there exist $0 < t_1 \leq t_2 \leq 1$ such that:

$$x(t) = \begin{cases} 1 & t \in [0, t_1), \\ x & t \in [t_1, t_2], \\ 0 & t > t_2, \end{cases}$$

where

$$x := \frac{t_1 (v_0^{t_1} - c(t_1))}{t_2 c(t_2) - (t_2 - t_1) v_{t_1}^{t_2} - t_1 c(t_1)}, \quad (6)$$

and $v_0^{t_1}$, $v_{t_1}^{t_2}$ are the conditional expectations of the buyer's value over the relevant intervals, namely

$$v_0^{t_1} := \frac{1}{t_1} \int_0^{t_1} v(t) dt, \quad v_{t_1}^{t_2} := \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} v(t) dt.$$

As can be verified, the threshold t_1 (resp., t_2) minimizes (resp., maximizes) the ratio

$$\frac{\int_{t_1}^{t_2} (v(t) - c(t)) dt}{\int_{t_1}^{t_2} t c'(t) dt},$$

given t_2 (resp., t_1). The numerator measures the gains from trade with the types in the interval $[t_1, t_2]$, while the denominator measures the information rents of the seller's types in that interval.¹² Indeed, if the buyer were to trade with, and only with, the seller's types $[0, t]$, his expected gains would be at most

$$Y(t) := \int_0^t (v(s) - c(t)) ds = \int_0^t (v(s) - c(s) - s c'(s)) ds, \quad (7)$$

a function that plays an important role in Samuelson and Myerson's analysis, as in ours.

3.2 The Veto-Incentive Compatible Program

Recall that the veto-incentive compatible program is obtained by adding to the full commitment program the requirement that, for any outstanding offer, the buyer's payoff is always non-negative, conditional on the outstanding offer, given his updated beliefs. At first sight, these

¹²To see this, note that, from the formula for Y given by (7), $\int_{t_1}^{t_2} s c'(s) ds$ is the difference between the gains from trade and the buyer's additional profit accruing from the types $[t_1, t_2]$.

constraints appear rather intractable, since these are restrictions on the marginal distributions over offers derived from the joint distribution over types and offers that a mechanism defines. The main result of this subsection establishes that, in fact, these constraints can be formulated in terms of the probabilities of trade alone. Therefore, as in the full commitment problem, it is enough to consider functions x , rather than distributions defined by μ , to determine the payoff set, so that standard variational techniques can be applied.

We first characterize implementable allocations, and then achievable payoffs. The following proposition, proved in Section 4, characterizes the set of allocations that can be implemented in the veto-incentive compatible program. Recall that incentive compatibility and individual rationality are minimal requirements.

Proposition 1 *An incentive compatible, individually rational allocation (x, p) is implementable in the veto-incentive compatible program if and only if, for all $t \in T$,*

$$B(t) = \int_t^1 x(s)[v(s) - p(s)]ds \geq 0.$$

Equipped with Proposition 1, it is then straightforward to characterize the set of payoffs that can be achieved in the veto-incentive compatible program.

Theorem 2 *Suppose that $c(1) \geq \int_T v(t)dt$. In the veto-incentive compatible program:*

1. *The payoff set can be obtained, without loss of generality, by assuming that all local incentive compatibility constraints bind, and that the highest seller type's payoff is zero: $\pi^S(1) = 0$;*
2. *The payoff set is spanned by the set of non-increasing functions $x : T \rightarrow [0, 1]$ subject to, for all $t \in T$,*

$$\int_t^1 [x(s)(v(s) - c(s)) - \int_s^1 x(u)dc(u)]ds \geq 0, \tag{8}$$

given expected payments, for all $t \in T$,

$$\bar{p}(t) = x(t)c(t) + \int_t^1 x(s)dc(s).$$

Note that the constraint in the second part of the theorem is simply the requirement that $B(t) \geq 0$ for all $t \in T$, given the definition of \bar{p} . Theorem 2.2 states that any non-increasing function $x \in [0, 1]$ satisfying $B(t) \geq 0$ for all t (a constraint that only involves the function x) is part of an allocation that is implementable in the veto-incentive compatible program, along with the expected payments defined in the theorem, and that these allocations are a sufficient class to generate all the payoffs that can be achieved in this program. Because of the veto-incentive compatibility constraint, the mechanism that is constructed in the proof of this theorem is not, however, a mechanism with deterministic prices.

The constraints $B(t) \geq 0$ (as stated in Theorem 2.2 in terms of the probabilities $x(t)$ only) are linear (in x) as well. It follows that the payoff set can be once again determined by using continuous linear programming. There is, however, one difficulty that is common to incentive problems with hidden characteristics and a continuum of types, namely the requirement that the function x be non-increasing. Fortunately, tools exist for such constraints. See, in particular, Hellwig (2009). What is the structure of the solution for boundary points of the payoff set? It depends, of course, on the specific boundary point and the underlying functions c and v . Note that, by differentiating (8) twice, we obtain that the probability x must satisfy the ordinary differential equation

$$x'(t)(v(t) - c(t)) + x(t)v'(t) = 0$$

on any such interval. The problem then reduces to identifying this finite partition. Indeed, examples can be constructed for which B is identically zero over some interval, and therefore, the allocation need not be a step function, nor the payoff set a convex polygon (the set of extreme points need not be finite).

It is an easy consequence of this theorem that the payoff vector maximizing the buyer's payoff in the veto-incentive compatible program coincides with the payoff vector that maximizes the

buyer's payoff in the full commitment program.¹³ The seller's highest payoff is either equal to, or smaller than the corresponding payoff in the full commitment program. Sufficient conditions for equality will be provided in the next section.

3.3 Bargaining Game

We finally consider the bargaining game. Clearly, for any history, given any outstanding offer that is accepted with positive probability, sequential rationality requires that the buyer's conditional payoff from accepting it must be non-negative. Therefore, the *ex ante* payoffs that can be achieved via bargaining must form a subset of the payoff set of the veto-incentive compatible program. But bargaining imposes additional constraints. For instance, since v is non-decreasing, it is common knowledge that the object is worth at least $v(0)$ to the buyer. Therefore, the seller of type t can secure a payoff of $v(0) - c(t)$, since he can keep on making this offer. (The formal argument is standard and omitted. See, for instance, Fudenberg, Levine and Tirole (1985), Lemma 2, which establishes that no lower offer is ever submitted in equilibrium, so that any such offer is necessarily accepted.) It is worth pointing out here that, if (x, p) is incentive compatible, then $\pi^S(0) \geq v(0) - c(0)$ implies that $\pi^S(t) \geq v(0) - c(t)$ for all $t \geq 0$, so that the aforementioned requirement reduces to $\pi^S(0) \geq v(0) - c(0)$. Since this provides a lower bound on the seller's payoff, we may think of this as the seller's *reservation payoff* in the bargaining game, a strengthening of individual rationality. Note that the most efficient mechanism in the veto IC program automatically satisfies the reservation payoff constraint.

One might wonder whether bargaining imposes additional restrictions on achievable payoffs. The main result of this section states that this is not the case, at least when $\delta \rightarrow 1$. Before stating this result, note that, with any equilibrium σ and for each type t , one can associate a quantity $x(t)$, namely the discounted total probability with which trade occurs under σ given t ,

¹³In fact, this follows from Proposition 1 in Samuelson (1984), as he shows that the buyer's favorite outcome is a take-it-or-leave-it offer, so that veto-incentive compatibility does not bind at this allocation.

$$x(t) = \mathbb{E}_\sigma \left[\sum_n \delta^{n-1} \mathbf{1}_{\sigma_n^B(h^{n-1}, p_n)=1} \right],$$

where $\mathbf{1}_A$ is the indicator function of the event A . Similarly, given some equilibrium σ , we let $\bar{p}(t) \in \mathbb{R}$ denote the expected discounted payment received by type t in this equilibrium. References to local incentive compatibility, or individual rationality, can be understood in terms of the pair (x, \bar{p}) . Recall that $\underline{\Pi} := \liminf_{\delta \rightarrow 1} \Pi(\delta)$ and $\overline{\Pi} := \limsup_{\delta \rightarrow 1} \Pi(\delta)$.

Theorem 3 *Suppose that $c(1) \geq \int_T v(t) dt$. Then $\underline{\Pi} = \overline{\Pi} =: \Pi$. Further, this set of payoffs is equal to the set of payoffs that can be achieved by veto-incentive compatible allocations for which*

$$\pi^S(0) \geq v(0) - c(0).$$

This result establishes that the only additional constraint on payoffs imposed by the bargaining game is that the lowest seller's type must secure his reservation payoff. In terms of efficiency, for instance, this theorem implies that there is no difference between the best outcome in the bargaining game and the veto-incentive compatible mechanism.

However, it is not true that any individually rational, incentive compatible allocation satisfying veto-incentive compatibility, and which gives his reservation payoff to the seller, can be necessarily implemented in the bargaining game. In Section 5.3, we provide an example of such an allocation, and explain why it cannot be implemented. For any such allocation, our result implies that there exists a payoff-equivalent allocation (in terms of *ex ante* payoffs for the seller and the buyer) that can be implemented. Therefore, bargaining imposes restrictions on implementable allocations that go beyond veto-incentive compatibility (and the restriction imposed by the security payoff), but not on payoffs.

Which constraints bind depends on the vertex of the set Π that is considered. On the upper boundary of this set, it can be assumed, without loss of generality, that all local incentive compatibility constraints are binding, and that the highest type's payoff of the seller trading with positive probability is zero: $\pi^S(1) = 0$; on the other hand, for those vertices that minimize some

convex combination of the seller’s and buyer’s payoff, the incentive compatibility constraints bind “downward,” that is, for all $t \in T$,

$$\pi^S(t) = \lim_{s \uparrow t} \pi^S(s | t),$$

with the boundary condition that the trading price of the highest seller’s type t is given by the minimum of $v(t)$ and either $\lim_{s \downarrow t} c(s)$, $t < 1$, or $c(1)$ if $t = 1$.

Given that the bargaining game imposes only one additional linear constraint to the veto-incentive compatible program, it can be analyzed via linear programming as well. Depending on c and v , this additional constraint can create a discontinuity (*i.e.*, a step) in the function x which has no counterpart in the previous (veto-incentive compatible) program, and arises before the first binding constraint $B(t) = 0$. Notice also that the constraint that $\pi^S(0) \geq v(0) - c(0)$ implies that the seller secures the *ex ante* payoff $\mathbb{E}[[v(0) - c(t)]^+]$ (because, as already mentioned, it implies that $\pi^S(t) \geq v(0) - c(t)$ for all t). However, the two requirements are not equivalent, as the example in the next subsection illustrates.

The proof of Theorem 3 is sketched in the next section. In doing so, we show that the payoff vector maximizing the seller’s payoff, which is also the efficient payoff vector in this set, coincides with the payoff vector maximizing the seller’s payoff in the veto-incentive compatible program. That is, as far as efficiency is concerned, bargaining imposes no constraint beyond veto-incentive compatibility. In all three programs, the *ex ante* buyer’s payoff is zero in any efficient allocation.

The proof is by construction. This requires us to specify beliefs after out-of-equilibrium offers. While sequential equilibrium is not well-defined in this game (the action space being infinite), our equilibrium can be made sequential by restricting this action set to a sufficiently rich but finite set of values. In this sense, our choice of off-path beliefs, while dictated by convenience, is not particularly fragile. Refinements just do not have much bite in an environment as rich as ours, and even Markov equilibrium does not appear to narrow down the payoff multiplicity. (A proof of this claim can be found in our online Appendix H, in which we prove that, at least in

the case in which there are finitely many types, the equilibrium strategies used in the proofs can be modified so as to be Markov.)

3.4 Examples and Economic Implications

To illustrate the results, we consider here an example with three types.

Example 4 *The functions v and c are step functions with three steps, and the two discontinuities occur for both functions at $t = 1/3$ and $2/3$. To simplify, we refer to those three types as 1, 2, and 3, assumed to be equally likely. Values and costs are given by*

$$(c_1, c_2, c_3) = (0, 4, 9), \text{ and } (v_1, v_2, v_3) = (2, 5, 12),$$

so that a higher index means a higher value, but also a higher cost. The left panel of Figure 1 represents the three payoff sets. The largest area is the set of payoffs in the full commitment case, and the smallest payoff set is the equilibrium payoff set in the bargaining game as $\delta \rightarrow 1$. In between lies the set of veto-incentive compatible payoff vectors. By changing only one parameter, namely, by increasing v_2 from 5 to 10, the payoff sets change considerably. See right panel. The two points $(440/1323, 20/63)$ on the left, and $(56/243, 2/9)$ on the right, represent the unique equilibrium payoff vectors in the bargaining game in which the (uninformed) buyer makes the offers in every period, as characterized in Deneckere and Liang (2006; “DL” in the figure) for $\delta \rightarrow 1$. As is clear, this (buyer-proposing) equilibrium payoff need not lie in the set of (seller-proposing) equilibrium payoff vectors, although it achieves a lower surplus than the maximum joint surplus when the seller makes offers. This is no coincidence, see below.

This example illustrates several points that hold more generally. First, as mentioned, the buyer’s highest payoff coincides in the veto-incentive compatible and the full commitment programs, but it might be lower in the equilibrium of the bargaining game. The seller’s highest payoff coincides in the bargaining game and the veto-incentive compatible program. This high-




Figure 1: Full commitment, Veto-Incentive Compatible, and Limiting Equilibrium Payoff Sets.

est payoff, however, might fall short of the highest payoff in the commitment program.¹⁴

When is (constrained) efficiency possible under bargaining, *i.e.*, when is veto-incentive compatibility consistent with efficiency? Obviously, this is trivially the case if the optimal allocation under full commitment is such that no seller's type trades with interior probability. If some seller's types do trade with interior probability, sufficient conditions can be given in terms of the buyer's gain function Y (see (7)). Let \underline{t} denote the smallest local maximizer of Y (which exists because $Y(0) = 0$ and $Y'(0) > 0$, yet $Y(1) < 0$). Note that \underline{t} solves $v(\underline{t}) - c(\underline{t}) = \underline{t}c'(\underline{t})$ (assuming differentiability at this point for the sake of this discussion). Let also \bar{t} denote the smallest strictly positive root of Y . We show in the online appendix (Appendix D) that efficiency

¹⁴Note also that, as is clear from the left panel, the restriction on achievable payoffs imposed by the lowest seller's type reservation payoff is not equivalent to the restriction that the seller obtains the *ex ante* payoff $\mathbb{E}[[v(0) - c(t)]^+] = 2/3$. Consider the vertex that minimizes the seller's payoff, subject to the buyer's payoff being zero. The requirement that the seller's lowest type gets at least $v(0) - c(0)$ drives the seller's *ex ante* payoff up to $17/18 > 2/3$. In this example, driving the seller's *ex ante* payoff down to $\mathbb{E}[[v(0) - c(t)]^+]$ is only possible in some equilibrium for high enough values of the buyer's payoff.

is attainable in bargaining if

$$\forall t \geq \bar{t} : \int_{\underline{t}}^t (v(s) - c(t)) ds \geq 0. \quad (9)$$

Obviously, as our example above shows (left panel), it is not always true that efficiency can be achieved. Note that the condition becomes easier to satisfy as gains from trade ($v(t) - c(t)$) increase, and information rents ($tc'(t)$) decrease (both \underline{t} and \bar{t} then increase). We summarize this discussion as follows.

Constrained efficiency can be achieved by bargaining as $\delta \rightarrow 1$ (even when the first-best outcome cannot) if gains from trade are high, or information rents low enough.

As an illustration, we consider Samuelson's Example 1, in which $c(t) = t$, and $v(t) = kt + \Delta$, where $k, \Delta \geq 0$. (This example subsumes both Akerlof's linear example ($\Delta = 0$) and the uniform additive example ($k = 1$) of Myerson, 1985.) See Figure 2 and Appendix G for details on Condition (9) in this example. Gains from trade require that $\Delta > 1 - k$, an area to which we can restrict attention. When $k \geq 2$ or $\Delta \geq 1 - k/2$ (Area *B*), the first-best is implementable in the veto-incentive compatible (and *a fortiori* in the full commitment) program. In Area *A*, when $\Delta \geq \frac{4}{4-k} - k$, Condition (9) is satisfied. In particular, in Area *A*, the first-best is not implementable under full commitment, but imposing veto-incentive compatibility comes at no additional cost. In the remaining area (for $\Delta \in [1 - k, 4/(4 - k) - k)$), Condition (9) is not satisfied, yet the conclusion remains valid: veto-incentive compatibility comes "for free." This suggests that veto-incentive compatibility is a relatively weak constraint (but also that the linear-additive example is somewhat misleading, cf. Example 4). Furthermore, it is more likely to be satisfied when gains from trade are high (as measured either by the boundary computed according to (9), or by the locus $\Delta = 1 - k$ that defines whether veto-incentive compatibility can be satisfied).

Because bargaining can achieve the same degree of efficiency as any (incentive compatible, individually rational) mechanism that satisfies veto-incentive compatibility, it follows that market




Figure 2: Sufficient conditions in Samuelson's Example 1.

institutions may only improve upon bargaining if they constrain the buyer in a way that weakens the veto-incentive compatibility constraint. This is demanding, but not impossible. For instance, screening contracts by the uninformed party (here, the buyer), as in Rothschild and Stiglitz (1976), dispense with the requirement of veto-incentive compatibility: the uninformed party offers (and commits) to a menu of price and quantity pairs, and the informed party chooses from them. This is not as demanding in terms of commitment as full commitment, but the difference is small (see Mylovanov 2008). At any rate, there is little to gain from trading institutions that would be less restrictive. Communication, for instance, does not expand the set of equilibrium outcomes. (Formally, the set of allocations that are achieved by communication equilibria is the same as those achieved by perfect Bayesian equilibria in the bargaining game, as $\delta \rightarrow 1$). Fortunately, as discussed, circumstances in which veto-incentive compatibility does not reduce efficiency are common, and in those circumstances, as little commitment as called for by bargaining suffices.

How do equilibrium outcomes in bargaining compare with the unique equilibrium outcome derived by Deneckere and Liang, when the buyer makes all the offers? In our two examples, the seller does worse in the latter equilibrium outcome than in any equilibrium outcome of our game. However, it is easy to construct examples in which this is not the case. In fact, the following can

be shown (details available upon request).

Lemma 1

- i. The allocation from the unique limit equilibrium outcome of the game in which the buyer makes all the offers is an equilibrium allocation in the game in which the seller makes all the offers if and only if it gives the lowest seller's type his reservation payoff (i.e., $v(0) - c(0)$), provided that the discount factor is sufficiently close to one.*
- ii. For δ close enough to one, the game in which the seller makes all the offers admits an equilibrium outcome that is strictly more efficient than the limit equilibrium outcome of the game in which the buyer makes all the offers.*

The first statement should come as no surprise given that the allocation that results from the bargaining game in which the buyer makes all the offers must be veto-incentive compatible. This follows from the “skimming” property in bargaining: because, from any history onward, the remaining seller’s types are all types above some threshold z_n , and because the buyer’s continuation payoff must be non-negative, it must be that $B(z_n) \geq 0$.¹⁵

The second statement is immediately implied by the first, given that the buyer secures a strictly positive payoff when he makes the offers, yet within the set of veto-incentive compatible allocations, efficiency is maximized when the buyer gets zero profits.

Of course, this lemma compares the best equilibrium outcome in one game with the unique equilibrium outcome in the other. There might be equilibria in the game in which the seller makes all the offers that are more inefficient than the equilibrium outcome when the buyer makes offers. Rather surprisingly, our example illustrates that this need not be true, however. As is obvious from the right panel of Figure 1, efficiency might be necessarily higher when the seller makes all the offers. This makes apparent that having the seller make all the offers does not simply “expand” the set of equilibria.

¹⁵If z_n and z_{n+1} denote consecutive threshold types, the inequality $B(t) \geq 0$ for $t \in (z_n, z_{n+1})$ follows from the fact that the types in $[z_n, t]$ are the most unprofitable ones (for the buyer) above z_n .

4 Main Proofs

4.1 Proof of Proposition 1 and Theorem 2

The proof of Theorem 2 will be divided in several steps. First, we establish Proposition 1, which immediately implies Theorem 2.2, given Theorem 1. We will then show how this, along with some other observations, can be used to establish Theorem 2.1.

The proof of Proposition 1 is itself divided into three parts. First, we show that, given an allocation (x, p) , the condition that $B(t)$ be non-negative for all t is necessary for the allocation to be implementable in the veto-incentive compatible program. Second, we turn to sufficiency. We first show that the conditions are sufficient if the functions c and v are step functions. Then we show how, by appropriate limiting arguments, the result follows for any functions c and v satisfying the assumptions of the model. For the sake of concision, we relegate the second and third part to an appendix (online Appendix A); but the gist of the idea can be conveyed in the simple case in which cost and value functions are step functions with two steps only, and that this is what is done in Lemma 3 below.

4.1.1 Proof of Proposition 1

First comes necessity.

Lemma 2 *If (x, p) is an allocation that is implementable in the veto-incentive compatible program, then, for all $t \in T$,*

$$B(t) = \int_t^1 (x(s)(v(s) - p(s))ds \geq 0.$$

Proof. Fix an allocation (x, p) that is implementable in the veto-incentive compatible program, and let μ denote the corresponding mechanism. Observe that, for all $t \in T$,

$$\begin{aligned}
\int_t^1 x(s)p(s)ds &= \int_t^1 \int_{\mathbb{R}_+} p\mu(s)[1, dp]ds \\
&\leq \int_t^1 \int_{\mathbb{R}_+} \left(\int_T v(u)d\nu(u | p) \right) \mu(s)[1, dp]ds \\
&\leq \int_t^1 \int_{\mathbb{R}_+} \frac{\int_{u \geq t} v(u)d\nu(u | p)}{\nu([t, 1] | p)} \mu(s)[1, dp]ds \\
&= \int_t^1 x(s)v(s)ds.
\end{aligned}$$

The first equality follows from the definition of the function p (see (2)). The first inequality is implied by veto-incentive compatibility; the second follows from the monotonicity of v ;¹⁶ the last equality, from the law of iterated expectations. This establishes the claim. \blacksquare

We now show sufficiency in the special case in which c and v are step functions with only one jump. As mentioned, the complete proof is in online Appendix A.

Lemma 3 *Suppose that c and v are step functions with a unique discontinuity point at $\hat{t} \in (0, 1)$. Consider the allocation (x, p) defined by*

$$x(t) = \begin{cases} x_1 & \text{if } t < \hat{t} \\ x_2 & \text{if } t \geq \hat{t} \end{cases} \quad p(t) = \begin{cases} p_1 & \text{if } t < \hat{t}, \\ p_2 & \text{if } t \geq \hat{t}. \end{cases}$$

Suppose that (x, p) is implementable in the full commitment program, and for every $t \in T$,

$$B(t) = \int_t^1 (x(s)(v(s) - p(s))ds \geq 0.$$

Then (x, p) is also implementable in the veto-incentive compatible program.

¹⁶That is, with the understanding that the integrand is zero when $\nu([t, 1] | p) = 0$.

Let v_1 and $v_2 \geq v_1$ denote the two values that the function v takes. The hypothesis that $B(t) \geq 0$ for every $t \in T$ implies that $v_2 \geq p_2$ and

$$\hat{t}x_1(v_1 - p_1) + (1 - \hat{t})x_2(v_2 - p_2) \geq 0.$$

Incentive compatibility implies that $x_1 \geq x_2$ and $p_1 \leq p_2$. It follows that there exists $z \in \left[0, \frac{x_2}{x_1}\right]$ such that

$$zx_1p_1 + (x_2 - zx_1)v_2 = x_2p_2. \quad (10)$$

To see this, notice that for $z = 0$, the left-hand side reduces to x_2v_2 , which is (weakly) larger than the right-hand side. On the other hand, for $z = \frac{x_2}{x_1}$, the left-hand side reduces to x_2p_1 , which is (weakly) smaller than the right-hand side.

Consider now the following random mechanism. When the seller announces a low type (*i.e.*, a type smaller than \hat{t}), the buyer receives the offer p_1 with probability x_1 (with probability $1 - x_1$ no offer is made and there is no trade). When the seller announces a high type, the buyer receives the offer p_1 with probability zx_1 and the offer v_2 with probability $x_2 - zx_1$ (again, there is no trade with the remaining probability).

It is immediate to check that the above mechanism implements the allocation (x, p) . Consequently, it satisfies the seller's incentive compatibility and individual rationality constraints. We now show that the mechanism also satisfies the buyer's veto incentive compatibility constraints. Notice that the buyer receives the offer v_2 only if the seller's type is high. Clearly, the buyer is willing to accept that offer. Finally, suppose that the buyer receives the offer p_1 . His expected payoff (conditional on p_1) is equal to

$$\frac{1}{\hat{t}x_1 + (1 - \hat{t})zx_1} [\hat{t}x_1(v_1 - p_1) + (1 - \hat{t})zx_1(v_2 - p_1)].$$

We multiply the above expression by $\hat{t}x_1 + (1 - \hat{t})zx_1$ and add and subtract $(1 - \hat{t})x_2(v_2 - p_2)$ to obtain

$$\begin{aligned} \hat{t}x_1(v_1 - p_1) + (1 - \hat{t})x_2(v_2 - p_2) - (1 - \hat{t})x_2(v_2 - p_2) + (1 - \hat{t})zx_1(v_2 - p_1) = \\ \hat{t}x_1(v_1 - p_1) + (1 - \hat{t})x_2(v_2 - p_2) \geq 0, \end{aligned}$$

where the equality follows from the definition of z in (10). We conclude that the buyer is also willing to accept the offer p_1 and that the allocation (x, p) is implementable in the veto-incentive compatible program.

4.1.2 Proof of Theorem 2.1

To establish Theorem 2.1, it remains to show that the payoff set of the veto-incentive compatible program can be obtained by assuming that:

1. all local incentive compatibility constraints are binding;
2. the highest type of the seller that trades with positive probability has a zero payoff;
3. the lowest type of the seller trades with probability 1, that is $x(0) = 1$.

Let us refer to the resulting payoff set as Π^V . Note that this set is compact and convex. These three claims are established by considering the boundary of Π^V . Because both properties are preserved under convex combinations, the result follows for the entire set. The details are somewhat tedious and also relegated to Appendix A.

4.2 Proof of Theorem 3

This theorem is proved by considering the set of extreme points of the payoff set, distinguishing them according to whether an extreme point lies to the “north-east,” “north-west,”

or “south-west” (*i.e.*, according to the signs of the weights on the seller’s and buyer’s payoff whose linear combination this extreme point maximizes). Arguments for one case require minor modifications for the other cases.¹⁷ For brevity, we only provide the complete proof for the case of positive weights, that is, we consider extreme points that lie on the Pareto-frontier.

The proof is divided into two steps. First, we show how allocations for which x is a step function satisfying some properties can be implemented as equilibria. Second, we show that every vertex of the equilibrium payoff set is the limit of a sequence of such allocations.

4.2.1 Regular Allocations

We first define a class of allocations (x, p) . Recall that, for $t_1 < t_2$, $v_{t_1}^{t_2} := \mathbb{E}[v(t) \mid t \in [t_1, t_2]]$.

Definition 1 *The allocation (x, p) is regular if there exist $0 = t_0 < t_1 < \dots < t_K \leq 1$, for some finite K , such that:*

1. *For some thresholds $1 = x_1 > \dots > x_K > 0$,*

$$x(t) = \begin{cases} x_k & \text{if } t \in [t_{k-1}, t_k), \quad k = 1, \dots, K, \\ 0 & \text{if } t \geq t_K; \end{cases}$$

2. *For some prices $v(0) \leq p_1 < \dots < p_K$,*

$$p(t) = \begin{cases} p_k & \text{if } t \in [t_{k-1}, t_k), \quad k = 1, \dots, K, \\ 0 & \text{if } t \geq t_K; \end{cases}$$

3. *For each $k = 1, \dots, K - 1$,*

$$x_k(p_k - c(t_{k,-})) = x_{k+1}(p_{k+1} - c(t_{k,-})),$$

where $t_{k,-} = \lim_{t \uparrow t_k} t$, and $x_K(p_K - c(t_{K,-})) \geq 0$;

¹⁷For instance, in the “south-west” region, the local incentive constraints are binding “downward,” and the definition of regular allocations must be modified accordingly.

4. Furthermore, we have $B(0) \geq 0, B(t_1) = \dots = B(t_{K-2}) = 0$, and $B(t_{K-1}) > 0$;

5. Finally, $v_{t_{K-2}}^{t_{K-1}} > c(t_{K-1}, -)$.

That is, a regular allocation is a step allocation such that local incentive compatibility constraints hold at each jump, the contribution to the buyer's payoff of each interval of types $[t_k, 1]$ is zero except for $k = 0, K - 1$, and positive for $t = 0, K - 1$ (strictly so for $t = K - 1$). Furthermore, the expected valuation of the buyer over the penultimate interval of types exceeds the cost of the seller's highest type in the previous interval, and the seller's lowest type must guarantee his security payoff.

Regular allocations are perfect candidates for equilibrium outcomes: one may think of each jump as defining an *interim* allocation (truncated according to the jumps) that satisfies incentive compatibility (for the seller) and individual rationality for the buyer at every step, as well as (by (1) and (2)) for the seller *ex ante*. In addition, regular allocations span a rich set of payoff vectors, as we will show that we can approximate every extreme point of the payoff set by a sequence of payoffs of regular allocations.

But a regular allocation need not be an equilibrium allocation in the discrete-time game, because of the indivisibilities that discrete periods introduce. This indivisibility becomes less and less problematic as $\delta \rightarrow 1$, and we show that we can choose (x^δ, p^δ) such that

$$\|(x^\delta, p^\delta)\| \rightarrow \|(x, p)\|,$$

uniformly in t , as $\delta \rightarrow 1$. The following lemma will be established in the next two subsections.

Lemma 4 *Fix a regular allocation (x, p) . There exists a sequence of $\sigma^\delta \in E(\delta)$ such that the corresponding sequence of allocations (x^δ, p^δ) converges to (x, p) as $\delta \rightarrow 1$, uniformly in t .*

In what follows, we consider regular allocations with $K > 2$ and $v_{t_{K-2}}^{t_{K-1}} > v(0)$ (notice that the last condition is automatically satisfied by regular allocations with $K > 3$). This simplifies

the notation since we are able to show that for δ sufficiently large the regular allocation (x, p) is an equilibrium allocation (*i.e.*, we have exact implementation). The proofs for the remaining cases are very similar, but require additional notation (in such cases we only have approximate implementation). We first establish some properties that regular allocations satisfy and then present the equilibrium of the bargaining game.

4.2.2 Properties of Regular Allocations

It follows from $B(t_{K-2}) = 0$, $B(t_{K-1}) > 0$ and $p_K > p_{K-1}$, that $v_{t_{K-2}}^{t_{K-1}} < p_{K-1} < v_{t_{K-2}}^{t_K}$. Thus, there exists a unique $\beta \in (0, 1)$ such that

$$p_{K-1} = \beta v_{t_{K-2}}^{t_K} + (1 - \beta) v_{t_{K-2}}^{t_{K-1}}. \quad (11)$$

Using $B(t_{K-2}) = 0$, $v_{t_{K-2}}^{t_{K-1}} > c(t_{K-1}, -)$ and the incentive compatibility constraint

$$x_{K-1}(p_{K-1} - c(t_{K-1}, -)) = x_K(p_K - c(t_{K-1}, -)),$$

it is easy to show that β also satisfies

$$x_K - x_{K-1}\beta = x_{K-1}(1 - \beta) \left(\frac{v_{t_{K-2}}^{t_{K-1}} - c(t_{K-1}, -)}{v_{t_{K-1}}^{t_K} - c(t_{K-1}, -)} \right) > 0,$$

and

$$x_{K-1}\beta v_{t_{K-2}}^{t_K} + (x_K - x_{K-1}\beta) v_{t_{K-1}}^{t_K} = x_K p_K$$

As the last step in the proof of Lemma 4, we show that the allocation (x, p) can be implemented in the bargaining game when the discount factor δ is sufficiently large.

4.2.3 The equilibrium σ^δ of the bargaining game

First, we describe the players' on-path behavior. Then we turn to the off-path behavior.

The behavior of the types in the intervals $[t_0, t_1)$ and $[t_K, 1]$ is very simple. In the first period of the game, the seller's types in $[t_0, t_1)$ make the offer p_1 and the buyer accepts it. In every period $n = 1, 2, \dots$, the types in $[t_K, 1]$ make the losing offer $v(1)$ and the buyer rejects it.

Consider now the types in the interval $[t_{k-1}, t_k)$, $k = 2, \dots, K-2$. In each period $n = 1, 2, \dots$, they make the offer p_k . In each period, the buyer accepts the offer p_k with probability $\psi_k = \frac{x_k(1-\delta)}{1-x_k\delta}$. Therefore, the discounted probability that the offer p_k is accepted is equal to

$$\psi_k + \delta(1 - \psi_k)\psi_k + \delta^2(1 - \psi_k)^2\psi_k + \dots = x_k.$$

The remaining types to consider are those in the last two intervals. Let \hat{n} denote the integer that satisfies¹⁸

$$\delta^{\hat{n}+2} \leq x_{K-1} < \delta^{\hat{n}+1}.$$

Also, let $\hat{\beta} \in (0, 1)$ be such that $\delta^{\hat{n}}\hat{\beta} = x_{K-1}\beta$.

The types in $[t_{K-2}, t_{K-1})$ and the types in $[t_{K-1}, t_K)$ adopt the same behavior in the first \hat{n} periods of the game. In particular, in period $n = 1, \dots, \hat{n} - 1$ they all make the losing offer $v(1)$ which the buyer rejects. In period \hat{n} , these types make the offer $v_{t_{K-2}}^{t_K}$. The buyer accepts this offer with probability $\hat{\beta}$.

If the buyer rejects the offer $v_{t_{K-2}}^{t_K}$ in period \hat{n} , then the types in the two intervals behave differently. In each period $\hat{n} + 1, \hat{n} + 2, \dots$, the types in $[t_{K-2}, t_{K-1})$ make the offer $v_{t_{K-2}}^{t_{K-1}}$. In each period, the buyer accepts this offer with probability ψ_{K-1} . This probability is such that

$$x_{K-1} = \delta^{\hat{n}} \left(\hat{\beta} + (1 - \hat{\beta}) \delta \left(\frac{\psi_{K-1}}{1 - \delta(1 - \psi_{K-1})} \right) \right).$$

¹⁸Notice that \hat{n} is well defined for δ sufficiently close to one.

Notice that $\psi_{K-1} \in (0, 1)$ since $\delta^{\hat{n}+1} > x_{K-1}$. Furthermore

$$\lim_{\delta \rightarrow 1} \frac{\psi_{K-1}}{1 - \delta(1 - \psi_{K-1})} = 1. \quad (12)$$

Finally, in each period $\hat{n} + 1, \hat{n} + 2, \dots$, the types in $[t_{K-1}, t_K)$ make the offer $v_{t_{K-1}}^{t_K}$. The buyer accepts each offer with probability ψ_K , where $\psi_K \in [0, 1]$ is such that

$$x_K = \delta^{\hat{n}} \left(\hat{\beta} + (1 - \hat{\beta}) \delta \left(\frac{\psi_K}{1 - \delta(1 - \psi_K)} \right) \right).$$

Again, $\psi_K \in (0, 1)$ since $x_K > x_{K-1}\beta$ and $\delta^{\hat{n}+1} > x_{K-1} > x_K$.

Notice that if both players follow the behavior that we have just described, then each type $t \in [t_{k-1}, t_k)$, $k = 1, \dots, k$, trades the good with discounted probability equal to x_k and receives the discounted expected transfer $x_k p_k$.¹⁹ In other words, the players' behavior implements the regular allocation (x, p) .

To see that this behavior is part of an equilibrium, consider all possible deviations in turn. First, we analyze the buyer's deviations. Notice that there is only one deviation that is detectable and that, at the same time, does not end the game. This happens when the buyer rejects the offer p_1 in the first period. Then the types in $[t_0, t_1)$ keep making the same offer p_1 until the buyer accepts it. On the other hand, the buyer accepts the serious offer p_1 in the first period in which is made. Finally, suppose that the buyer deviates when he is supposed to randomize. Then, following this deviation (which is not detectable), he reverts to the behavior specified above.

Consider now the seller's deviations. The buyer accepts an off-path offer if and only if the offer is weakly smaller than $v(0)$. Following an off-path history of offers, type t of the seller offers $v(0)$ in every future period if $v(0) \geq c(t)$. Otherwise, type t offers $v(1)$ in every future period.

It is simple to verify that the strategy profile just described constitutes an equilibrium (or rather, that there exists a belief system along which this strategy profile is an equilibrium) when

¹⁹As mentioned above, the types above t_K do not trade the good and do not receive a transfer.

δ is sufficiently large. By construction, each type t prefers his own strategy to the strategy of type t' . Thus, no type $t \in [0, 1]$ has an incentive to mimic the equilibrium behavior of another type t' . Also, type t does not have any incentive to deviate and make an off-path offer strictly larger than $v(0)$ (the offer would be rejected) or strictly smaller than $v(0)$ (the type would be strictly better off by making the off-path offer $v(0)$). Finally, it follows from equality (12), from the condition $v_{t_{K-2}}^{t_{K-1}} > v(0)$, and from the incentive compatibility constraints that for δ sufficiently close to one, no type has an incentive to deviate and make the off-path offer $v(0)$.

Conditional on receiving the offer p_1 in the the first period, the buyer's expected payoff is weakly positive. Thus, he has an incentive to accept the offer. Conditional on receiving any other on-path offer, the buyer's expected payoff is zero. Therefore it is optimal to randomize.

The off-path behavior can be easily made sequentially rational by assuming that, after any deviation, the buyer assigns probability one to the event that the seller's type is $t = 0$. (This is a rather extreme belief revision, but it is convenient, and other possibilities would do just as well.)

4.2.4 Proof of Theorem 3, Conclusion

The previous subsections have shown that any regular allocation can be achieved as an equilibrium allocation in the bargaining game as $\delta \rightarrow 1$. Note that the set of equilibrium payoffs that can be achieved in the bargaining game is a subset of the set of payoffs spanned by the allocations described in Theorem 3, because the constraint $\pi^S(0) \geq v(0) - c(0)$ must hold, as explained. Also, equilibrium allocations must satisfy veto-incentive compatibility. Therefore, one direction of the Theorem 3 is obvious. The other direction will be established if we can show that every extreme point of the set of veto-incentive compatible payoffs giving the seller his security payoff can be approximated arbitrarily closely by regular allocations. This is the content of Lemma 5. Recall that, for brevity, we restrict ourselves here to the case of extreme points of the payoff set that lie on the Pareto-frontier. The following is proved in the online appendix (Appendix B). This Appendix also shows how the result extends to the case of a finite (rather than infinite)

horizon. The following lemma concludes the proof of Theorem 3.

Lemma 5 *For every extreme point (π^S, π^B) (on the north-east boundary) of the payoff set that can be achieved by veto-incentive compatible allocations for which $\pi^S(0) \geq v(0) - c(0)$, and every $\varepsilon > 0$, there exists a regular allocation whose payoff is within distance ε of (π^S, π^B) .*

5 Extensions

This section addresses three issues. First, we show that veto-incentive compatibility is implied by standard bargaining protocols (Section 5.1), as mentioned in the introduction. Second, we discuss the impact of imposing similar requirements on the seller’s side (Section 5.2). Modern commercial law emphasizes buyer’s rights. It is then natural to ask whether this is compatible with protection of the seller, and whether mechanisms can be found that are agreeable to both the buyer *and* the seller not only *ex ante*, but also *ex interim*.

Finally, in Section 5.3, we dwell on an important caveat. Our results characterize those *ex ante* payoffs that can be achieved under different levels of commitment. These characterizations do not carry over from payoffs to allocations, as we explain.

5.1 Bargaining Outcomes Satisfy Veto-Incentive Compatibility

Veto-incentive compatibility is not only sensible from a practical point of view, it can also be shown to be automatically satisfied by most bargaining protocols that appear in the literature. This is formally established here. We define bargaining games satisfying the following extensive forms: (i) Nature selects one party to make an offer at the first stage $n = 0$ (independently of the type). This party offers a price $p \in \mathbb{R}$, whereupon the other party decides to accept it or not. If the other party accepts and the seller’s type is t , then the buyer obtains a payoff $v(t) - p$, whereas the seller obtains a payoff $p - c(t)$. If the offer is rejected, Nature ends the game with probability $\theta(h^1)$, where h^1 is the public history. If the game does not end, nature

determines that the buyer makes an offer with probability $\chi(h^1)$ and the seller makes an offer with probability $1 - \chi(h^1)$. The game proceeds accordingly. The non-increasing sequence $\{\delta_n\}_{n \geq 0}$ defines the common discount factor of period n (as evaluated from period 0, *e.g.*, in the case of geometric discounting, $\delta_n = \delta^n$). We normalize δ_0 to 1 and assume continuity at infinity, that is, $\lim_{n \rightarrow \infty} \delta_n = 0$. Our protocol encompasses alternate-offers protocols, protocols in which the buyer makes all the offers, as well as finite-horizon protocols (set $\delta_n = 0$ for all $n > T$). Any such extensive form, alongside the incomplete information and the utility function, defines a multi-stage game of observed actions. We define perfect Bayesian equilibria of such a game and the corresponding allocation as we did in Section 3.3. Proposition 2 establishes that any equilibrium of the bargaining game satisfies veto-incentive compatibility. (See Online Appendix F for a proof.)

Proposition 2 *Any perfect Bayesian equilibrium allocation satisfies veto-incentive compatibility.*

5.2 Limited Commitment on the Seller’s Side

Veto-incentive compatibility weakens the commitment assumption made in the full commitment program on the buyer’s side. As discussed, this is a relaxation that is relevant for many actual market institutions. Furthermore, our characterization of the equilibrium payoffs in the bargaining game suggests that this is the “right” relaxation, namely, the absence of commitment on either side, as captured by the bargaining game, appears to impose no further constraints on achievable payoffs, aside from the security payoff that the seller must secure.²⁰

It is then natural to ask whether one could derive results that mirror those of Section 3.2 in which the seller’s commitment, instead of, or in addition to, the buyer’s commitment is relaxed. While we shall not attempt to obtain a characterization for each possible case, we discuss here

²⁰Of course, in bargaining, the seller is not formally allowed to withdraw an offer that he makes, but why would he? Acceptance by the buyer reveals no information, so a seller that anticipates withdrawing an offer might as well not submit it.

the relationship between the different sets of allocations and payoffs. As we shall see, limited commitment on the seller's side is arguably less of a problem than on the buyer's side.

Unlike the buyer, the seller gets an opportunity to influence the terms at which the trade would take place. Therefore, there are two possible ways of modeling the absence of commitment on the seller's side. A mechanism is *ex post* individually rational for the seller if the price p that is offered to the buyer is always higher than the cost of the seller's reported type t :

$$\forall t \in T : \int_{[0, c(t))} \mu(t)[1, dp] = 0.$$

This guarantees that the seller never loses from the mechanism, but it does not give him the authority to actually prevent the trade. Alternatively, we might endow the seller with the ability to block the trade given the realized price. This notion, in line with Forges' original definition of veto-incentive compatibility, is more demanding than *ex post* individual rationality: the ability to block the trade affects the seller's incentives to report his type truthfully, as the payoff from making a given report must include the option value from blocking the trade if the realized price happens to be below the seller's actual cost. To be more formal, we re-define the payoff of the type t seller that reports s , from a given mechanism μ , as

$$\hat{\pi}^S(s | t) = \int_{\mathbb{R}_+} \mathbf{1}_{\{p \geq c(t)\}} (p - c(t)) \mu(s)[1, dp].$$

A mechanism is *seller veto-incentive compatible* if it is incentive compatible given the payoff $\hat{\pi}$, and the allocation (x, p) is implementable in the seller's veto-incentive compatible program if there is a mechanism that is seller veto-incentive compatible and induces the allocation (x, p) , according to (1)–(2), taking into account that trade does not take place for prices below $c(t)$. To distinguish this notion from veto-incentive compatibility as defined in Section 2, the latter will now be referred to as buyer veto-incentive compatibility.

Does seller veto-incentive compatibility, or even *ex post* individual rationality restrict the set of implementable allocations, or the set of achievable payoff vectors? In a nutshell, the answer is

no, as far as payoffs are concerned, and sometimes, as far as allocations are concerned, but only if it comes in addition to buyer veto-incentive compatibility. Formally:

Proposition 3

- i. The set of implementable allocations (and thus, of achievable payoff vectors) in the full commitment program remains unchanged if seller veto-incentive compatibility is imposed.*
- ii. The set of implementable allocations (and thus, of achievable payoff vectors) in the buyer veto-incentive compatible program remains unchanged if seller ex post individual rationality is imposed.*
- iii. The set of achievable payoff vectors in the buyer veto-incentive compatible program remains unchanged if seller veto-incentive compatibility is imposed.*

Because seller veto-incentive compatibility implies seller *ex post* individual rationality, we have omitted some relationships that follow from the proposition. For instance, from (i), it follows that seller *ex post* individual rationality does not restrict the set of implementable allocations in the full commitment program. Furthermore, all remaining inclusions are strict: that is, for some parameters, the set of implementable allocations in the buyer veto-incentive compatible program is strictly reduced if seller veto-incentive compatibility is imposed, and, as we know, the set of implementable allocations in the veto-incentive compatible program is strictly contained in the set of allocations of the full commitment program, for some parameters.

The proofs of the claims in Proposition 3, some of which follow arguments that are similar to the other proofs in the paper, are sketched in the online appendix (Appendix E).²¹ Additional details, as well as examples establishing the strict inequalities, are available from the authors.

²¹More precisely, we show in Appendix E that seller *ex post* individual rationality does not restrict the set of allocations that can be achieved in the (buyer) veto-incentive compatible program, and that, as far as payoffs are concerned, we can also impose seller veto-incentive compatibility. In both cases, attention is restricted to finite types, and the result in the general case follows by standard limiting arguments.

5.3 Payoffs vs. Allocations

Our characterization of veto-incentive compatibility and equilibrium outcomes in the bargaining game were cast in terms of the agents' expected payoffs, not in terms of allocations. For some important criteria, this makes no difference: the efficient payoff, the buyer's highest payoff, for instance, is implemented by a unique allocation. However, our result does not extend to allocations in general. *Not every incentive-compatible allocation whose payoffs satisfy the conditions of the characterization is implementable.* For instance, not every allocation that gives the seller's lowest type a profit $\pi^S(0) \geq v(0) - c(0)$ need be implementable. Suppose that there are three equiprobable types of seller (and buyer), and consider parameters such that the highest cost, c_3 , is lower than the expected value of the lower two values, $(v_1 + v_2)/2$. Further, consider an incentive compatible allocation in which the buyer's expected payoff is 0, the highest seller's type does not trade, but the second highest does; this seller's intermediate type gets a positive profit, and the seller's lowest type gets a payoff exceeding $v(0) - c(0)$; by our results, the resulting expected payoffs are equilibrium payoffs in the bargaining game when frictions are small.²²

Yet this specific allocation, which requires the seller's high type not to trade, cannot be implemented in the bargaining game. To see this, note that the buyer will never accept an offer that gives him a strictly negative payoff, and therefore, because the buyer's expected payoff is zero, it must be that his expected payoff is also zero after any history, conditional on any offer that is submitted with positive probability. By the martingale property of beliefs, there is a sequence of equilibrium offers along which the buyer's expected value, conditional on these offers, is non-decreasing, and therefore, at least as large as $(v_1 + v_2)/2 > c_3$. This sequence of offers must involve offers accepted with positive probability, for otherwise the seller's intermediate type would not be willing to follow it. By mimicking this sequence of offers, the seller's highest type guarantees a strictly positive profit, a contradiction.

²²Such an example is easy to find with a mathematical software: for instance, it occurs for the parameters $c_1 = 1, c_2 = 5970/2142, c_3 = 175/51$, and $v_1 = 134/65, v_2 = 2458/509, v_3 = 5$. The allocation is $x_1 = 1, x_2 = 1309475796/1359864155, x_3 = 0, p_1 = 926734382/271972831, p_2 = 898659860/271972831, p_3 = 0$.

References

- Akerlof, G. (1970). “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, **84**, 488–500.
- Ausubel, L.M. and R.J. Deneckere (1989). “A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information,” *Journal of Economic Theory*, **48**, 18–46.
- Deneckere, R. and M.-Y. Liang (2006). “Bargaining with Interdependent Values,” *Econometrica*, **74**, 1309–1364.
- Forges, F. (1994). “Posterior Efficiency,” *Games and Economic Behavior*, **6**, 238–261.
- Forges, F. (1999). “Ex post individually rational trading mechanisms,” in: Alkan, A., C. Aliprantis et N. Yannelis (eds.), *Current trends in Economics*, Springer-Verlag, 157–175.
- Fudenberg, D., D.K. Levine and J. Tirole (1985). “Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information,” in A. Roth (ed.), *Game Theoretic Models of Bargaining*, Cambridge University Press, 73–98.
- Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge, MA: MIT Press.
- Gresik, T. (1991a). “Ex ante efficient, ex post individually rational trade,” *Journal of Economic Theory*, **53**, 131–145.
- Gresik, T. (1991b). “Ex ante incentive efficient trading mechanisms without the private value restriction,” *Journal of Economic Theory*, **55**, 41–63.
- Hellwig, M.F. (2009). “A Maximum Principle for Control Problems with Monotonicity Constraints,” Preprint of the Max Planck Institute for Research on Collective Goods, Bonn, 2008/4.
- Matthews, S. and A. Postlewaite (1989). “Preplay communication in two-person sealed-bid double actions,” *Journal of Economic Theory*, **48**, 238–263.

- Myerson, R.B. (1981). "Optimal Auction Design," *Mathematics of Operations Research*, **6**, 58–73.
- Myerson, R.B. (1985). "Analysis of Two Bargaining Games with Incomplete Information," in A. Roth (ed.), *Game Theoretic Models of Bargaining*, Cambridge University Press, 115–147.
- Mylovanov, T. (2008). "Veto-Based Delegation," *Journal of Economic Theory*, **138**, 297–307.
- Riley, J.G. and R. Zeckhauser (1983). "Optimal Selling Strategies: When to Haggle, and When to Hold Firm," *Quarterly Journal of Economics*, **76**, 267–287.
- Rothschild, M. and J. Stiglitz (1976). "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, **90**, 629–649.
- Rubinstein, A. (1982). "Perfect Equilibrium in a Bargaining Model," *Econometrica*, **50**, 97–109.
- Samuelson, W. (1984). "Bargaining under Asymmetric Information," *Econometrica*, **52**, 995–1005.