A Mechanism-Design Approach to Speculative Trade*

Kfir Eliaz† and Rani Spiegler‡

November 6, 2005

Abstract

When agents hold non-common priors over an unverifiable state of nature which affects the outcome of their future actions, they have an incentive to bet on the outcome. We pose the following question: what are the limits on the agents' ability to realize gains from speculative bets when their prior belief is private information? We apply a "mechanism design" approach to this question, in the context of a pair of models: a principal-agent model in which the two parties bet on the agent's future action, and a market model in which traders bet on the future price. We characterize interim-efficient bets in these environments, and their implementability as a function of fundamentals. In general, implementability of interim-efficient bets diminishes as the costs of manipulating the bet's outcome become more uneven across states or agents.

*We thank Barton Lipman, Eric Maskin, Ady Pauzner, Wolfgang Pesendorfer, Ariel Rubinstein and especially Eddie Dekel, for helpful conversations and comments. Financial support from the US-Israel Binational Science Foundation, Grant No. 2002298 is gratefully acknowledged.

†Dept. of Economics, NYU. 269 Mercer St., New York, NY 10003, E-mail: kfir.eliaz@nyu.edu, URL: http://homepages.nyu.edu/~ke7

‡School of Economics, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: rani@post.tau.ac.il. URL: http://www.tau.ac.il/~rani

1 Introduction

One of the primary tasks of the mechanism-design literature has been to draw the limits that asymmetric information imposes on the ability to realize gains from trade. A milestone in this literature was the result due to Myerson and Satterthwaite (1983), which stated that in a natural class of bilateral-trade environments, there exists no mechanism that weakly implements efficient trade in Bayesian Nash equilibrium.

This result, like the rest of the literature that followed in its wake, focuses exclusively on trade that is motivated by differences in *tastes*. In principle, one could pose the same set of questions when the motivation for trade is differences in *beliefs*. What are the limits on the ability to realize gains from speculative trade, when the agents' beliefs are not common knowledge? What are the mechanisms that enable agents to realize these gains? The mechanism-design literature has ignored these questions.

The primary reason for this neglect seems to be the ubiquity of the common-prior assumption in economic modeling. As the no-trade theorems (e.g., Milgrom and Stokey (1982)) have shown, common priors coupled with standard solution concepts rule out speculative trade. In this paper, we focus on environments in which agents have different prior beliefs regarding a state of nature that may affect the outcome of their future actions. This creates a motive among the agents to bet on the future outcome. The agents' priors are private information, although the distribution from which the priors are drawn is common knowledge. This element of asymmetric information may inhibit speculative bets. We apply a *"mechanism-design approach"* in order to examine whether this barrier to speculative trade can be overcome, how these barriers are affected by the fundamentals of the environment, and what kind of mechanisms can be employed to realize the gains from speculative trade.

The observation that asymmetric information may act as a barrier to speculative trade, even when the common-prior assumption is relaxed, has at least two precedents in the literature. Morris (1994) provided necessary and sufficient conditions (in terms of the structure of the agents' beliefs) for no-trade results to persist in environments with heterogeneous priors. Chung and Ely (2005) study the design of auctions in an environment with non-trivial high-order beliefs. In particular, they allow for heterogenous priors, and show that incentive-compatibility constraints exclude mutually beneficial bets as part of the revenue-maximizing mechanism.

We demonstrate our mechanism-design approach to speculative trade with a pair of models. We first present a simple principal-agent model, in which the two parties speculate over which of a pair of actions the agent will take in a future decision problem.

The agent's payoff from actions depends on an unverifiable state of nature. Therefore, although the two parties hold different prior beliefs over the state, they cannot bet on the state but only on the agent's action.

Our first model is meant to convey the basic intuition underlying our approach to speculation over a manipulable outcome. We then apply this intuition to study speculative trade in an imperfectly competitive market where traders bet on the future market price. We focus on an economy with m identical sellers and n identical buyers, in which traders have conflicting prior beliefs over the future level of external demand. We assume that neither the level of external demand nor the traders' actions are verifiable. Therefore, the traders can given an expression to their heterogenous beliefs only by betting on the future market price.

In both models, the bets are on outcomes which are susceptible to manipulation by some, or all of the agents. This is an important feature of our model. It implies that in order for a bet to be sustainable, its stakes cannot exceed the cost of unilateral manipulation of its outcome. But this means that potential gains from speculative bets are bounded as well. An "optimal bet" in such an environment maximizes these gains (formally, the sum of the agents' interim expected utilities, where each agent's expected utility is calculated according to his own prior), subject to the constraint that none of the agents wish to manipulate its outcome.

Bounded bets could be generated by alternative assumptions, such as risk aversion or liquidity constraints. We find our method appealing for a number of reasons. First, there are many real-life situations in which agents can manipulate the outcomes they are betting on. In addition to the situations captured by our pair of models, one may also think of gambling on the outcomes of contests, where a major concern for gamblers is the possibility that a contestant might try to lose on purpose in order to win a bet. Second, from a methodological point of view, quasi-linear utility and unbounded transfers are standard equipment in the mechanism design literature. Therefore, it makes sense to retain it, as we apply the mechanism-design approach to speculative trade.

Most importantly, the bounds on the stakes of bets in our model are *endogenous*. Therefore, we are able to obtain novel insights as to how implementability of "optimal bets" depends on payoff-relevant details. Indeed, the main result in the principal-agent model links the implementability of the optimal bet to the structure of the agent's state-dependent utility function. When the ratio between the loss from taking the wrong action in one state and the loss from taking the wrong action in the other state becomes closer to one, the set of distributions over prior beliefs for which the optimal bet is implementable expands. When the ratio is exactly one, the optimal bet is implementable for any distribution over priors.

Likewise, implementability of the optimal bet in the market model depends on the ratio between the cost of upward price manipulation for buyers and for sellers. Greater buyer-seller asymmetry in this respect implies greater difficulties in implementing the optimal bet. Shrinking the gap between the buyers' and sellers' valuation of the traded good reduces this asymmetry and therefore expands the set of distributions over priors for which the optimal bet is implementable. Increasing the number of sellers has the same effect (regardless of the number of buyers), as long as $m \neq n$. The case of $m = n$ has a special status in this model: the optimal bet is implementable for every distribution of prior beliefs, using a natural, auction-like mechanism.

The technical basis for these results is a formal analogy to a more conventional mechanism-design model due to Cramton, Gibbons and Klemperer (1987) (henceforth, referred to as CGK). Their work extended the Myerson-Satterthwaite analysis to general initial ownership structures, namely "partnerships". The problem of implementing optimal bets turns out to be analogous to the problem of dissolving a partnership efficiently. We demonstrate that the analogy is not merely formal, but also provides an insight into the nature of the mechanism-design problem in the context of speculative trade.

2 Betting on an agent's future action

Consider an agent who faces a choice between two actions: a or b. His payoff from each option depends on the state of nature. There are two possible states. The agent's vNM utility function is u in one state and v in the other. With slight abuse of notation, we denote states by the utility functions that characterize them, given by the following table:

	a	b
u	A	C
v	D	B

where $A - C \geq 0$ and $B - D \geq 0$, with at least one strict inequality.

The agent privately learns the state of nature before making his decision. A period before the realization of the state, the agent and another player, referred to as a "*speculator*", hold different beliefs regarding the realization of the state. These are purely differences in *prior* opinions. Let θ_1 and θ_2 be the prior probability assigned to state u by the speculator and the agent, respectively.

Because the speculator and the agent have different priors, they would find it mutually beneficial to bet on the future state of nature. However, since the state is privately observed by the agent, such a bet is unenforceable. Instead, the players can bet on the agent's action, which is assumed to be verifiable. We refer to the period in which the state is realized and the action is taken as period 2. The period in which the bet is negotiated is referred to as period 1.

A *bet* is a function $t : \{a, b\} \to \mathbb{R}$ that specifies for every action a monetary transfer (possibly negative) from the agent to the speculator. If the players agree on a bet t, it affects the decision problem faced by the agent, such that the agent's utility from an action x is $u(x) - t(x)$ in state u and $v(x) - t(x)$ in state v, and the speculator's utility from the agent's action x is $t(x)$, regardless of the state. If no bet is signed, the agent faces the "bare" decision problem and the speculator receives nothing.

This example captures in a stylized way a number of real-life situations. For instance, the two players can be interpreted as a buyer and a seller. If no deal is signed between them in period 1, the buyer can purchase the good from an alternative supplier. In period 2, the buyer learns which of two varieties of the good is more suitable for him. In period 1, he lacks this information. A bet signed between the buyer and the seller is essentially an advance contract which forces the seller to provide in period 2 whichever variety the buyer demands, at a given price (providing either variety is costless for the seller). Alternatively, the agent can be interpreted as a central bank of a small economy, facing a decision whether or not to devalue the currency, depending on the rate of inflation. The speculator can be interpreted as a big trader in the exchange market, who intends to earn speculative gains, due to conflicting beliefs regarding the future inflation rate.

2.1 Interim-efficient bets

Consider a bet t, and suppose that both players expect that the agent's actions in states u and v will be x^u and x^v. Denote $x = (x^u, x^v)$. Then, the speculator's interim expected payoff from (x, t) is $\theta_1 \cdot t(x^u) + (1 - \theta_1) \cdot t(x^v)$, while the agent's is $\theta_2 \cdot [u(x^u) - t(x^u)] + (1 - \theta_2) \cdot [v(x^v) - t(x^v)]$. The term "*interim*" is fitting because it refers to the players' expected payoffs upon learning their prior, which is drawn from a common distribution. Note that the sum of the players' interim expected payoffs can be conveniently written as

$$\theta_2 \cdot u(x^u) + (1 - \theta_2) \cdot v(x^v) + (\theta_1 - \theta_2) \cdot [t(x^u) - t(x^v)] \tag{1}$$

If the agent could commit to play $x^u \neq x^v$, there is no upper bound on the stakes of the bet that the two players would want to sign: if $\theta_1 > \theta_2$, they would set $t(x^u) \gg t(x^v)$, and if $\theta_1 < \theta_2$, they would set $t(x^v) \gg t(x^u)$. However, because the agent cannot commit to his second-period action, the players must take into account his ability to manipulate the bet's outcome. For instance, suppose that the players agree on a bet that satisfies $t(b) - t(a) > \max(u, v)$. Then, regardless of the state, the agent will prefer to choose a, because the amount he saves in side payments outweighs the loss from taking the wrong action in the "bare" decision problem. But if the agent takes the same action in both states, the players cannot benefit from betting on the agent's action. *Thus, in order to be sustainable, the bet must provide the agent with incentives to take different actions in different states.*[1]

A pair (x, t) is *interim Pareto efficient* if it maximizes (1) subject to the constraints:

$$\begin{aligned} u(x^u) - t(x^u) &\geq u(x^v) - t(x^v) \\ v(x^v) - t(x^v) &\geq v(x^u) - t(x^u) \end{aligned}$$

which we call "second-period incentive compatibility" (SPIC) constraints. If (x, t) is a solution to this constrained optimization problem, we refer to t as an *interim-efficient bet*. We refer to the expression (1), evaluated at an interim-efficient pair (x, t), as the *first-best* (FB) surplus.[2] Interim efficiency is the relevant "optimality" criterion because it captures the players' motive to engage in speculative trade. For any (x, t) which is interim-inefficient, the players can find another pair (x', t') which both of them will prefer to (x, t), given their priors.

It follows from (1) that if $\theta_1 > \theta_2$, the players would want to set $t(x^u) - t(x^v)$ to be equal to the *upper* bound implied by the SPIC constraints, $u(x^u) - u(x^v)$. In contrast, if $\theta_1 < \theta_2$, they would want to set $t(x^u) - t(x^v)$ to be equal to its *lower* bound implied by the SPIC constraints, $v(x^u) - v(x^v)$. Turning to the determination of x, it is easy to see that both bounds on $t(x^u) - t(x^v)$ are relaxed to their utmost when $x^u = a$ and $x^v = b$. Thus, we have the following characterization.

[1] Can the upper bound on the bet's stakes be overcome by some general message game that the players could carry out in the second period? Even if the state is commonly known in period 2, the assumption that there are only two players and the restriction to budget-balanced transfers imply that it cannot. Without a third player or the ability to "burn money", a second-period mechanism is unable to punish players for submitting untruthful messages.

[2] There is a slight abuse of terminology here, as it is customary to consider the "first best" as the solution to an optimization problem, unconstrained by any incentive compatibility concerns. Because the expression (1) is unbounded if we ignore the SPIC, we consider the "first-best" to be the highest attainable surplus, taking into account the SPIC but ignoring the *first-period* incentive constraints that arise when the agent's type is unknown.

Remark 1 *A pair (x,t) is interim efficient if and only if the following two conditions hold:*

(i) x is ex-post efficient - i.e., $x^u = a$ and $x^v = b$.

(ii) t satisfies:

$$t(a) - t(b) = \begin{cases} A - C & if \quad \theta_1 > \theta_2 \\ D - B & if \quad \theta_1 < \theta_2 \end{cases}$$

Thus, the bounds on the stakes of the optimal bet are determined by the agent's cost of manipulating the bet's outcome in each state. In state u, he will lose $A - C = u(a) - u(b)$ if he switches from a to b. In state v, he will lose $B - D = v(b) - v(a)$ if he switches from b to a. The upper bound on $t(a) - t(b)$ is binding when the highest prior on u is assigned by the speculator, and the lower bound on $t(a) - t(b)$ is binding when the highest prior on u is assigned by the agent.

An important feature of the interim-efficient bet is that no matter what the state is, the cost to the agent of manipulating the outcome of this bet is $(A - C) + (B - D)$. To see this, suppose the state is u. Because the interim-efficient bet is ex-post efficient, the agent is expected to choose action a and pay $t(a)$ to the speculator. If, instead, the agent chooses b, then he forgoes a net payoff of $A - C$ and raises his payment to the speculator by $B - D$. Hence, the total cost of deviating from b to a is the sum of these two amounts. Now suppose the state is v, and the agent contemplates deviating from b to a. The agent would incur two types of costs from this deviation: first, he would forgo $B - D$, the value of b in state v, net of the value of a in this state, and second, he would raise his payment to the speculator by $A - C$. Once again, the total cost of manipulating the outcome is $(A - C) + (B - D)$.

Corollary 1 *The FB surplus is* $\max(\theta_1, \theta_2) \cdot (A - C) + \max(1 - \theta_1, 1 - \theta_2) \cdot (B - D)$.

The FB surplus has an interesting interpretation: it is as if we gave, in each state, the entire "bare" surplus from the agent's optimal action to the player who assigned the highest prior to that state.

2.2 Implementation

We now turn to the question of whether the FB can be implemented, when the players' priors are not common knowledge. Specifically, we assume that the speculator privately observes θ_1 and the agent privately observes θ_2, and it is common knowledge that both priors are independently drawn from a continuous *cdf* F with support $[0,1]$.

To see why privately known priors could act as a barrier to mutually beneficial speculative bets, suppose that $A = B = 100$ and $C = D = 0$. Consider a bet in the form of a lottery ticket which entitles its owner the right to receive 100 from the other player if and only if the agent's second-period action is a. Suppose that in order to obtain the ticket, one must pay a fixed price of 50. Interim efficiency requires us to allocate the ticket to the player with the highest prior on u. Consider the following naive allocation rule: each player reports his prior on u, and the highest-reporting player gets the ticket. Given this rule, players have an incentive to pretend to have extreme beliefs. When their prior exceeds $\frac{1}{2}$, they want to win the ticket and therefore bias their report upward. When their prior falls below $\frac{1}{2}$, they do not want the ticket and therefore bias their report downward. As a result, the interim efficient allocation fails to be implemented.

We consider the problem of implementing the FB via a *direct* mechanism. This means that the agent and the speculator play a two-period game, denoted Γ. In the first period, each player submits a report $\hat{\theta}_i \in [0,1]$ (interpreted as that player's stated prior on u), or chooses not to participate. If at least one player chooses the latter, the agent faces the "bare" decision problem. If both players choose to participate, every pair of reports $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \hat{\theta}_2)$ is assigned a transfer function $t(x \mid \hat{\boldsymbol{\theta}})$, which is disclosed to the agent. In period 2, after the state of nature is realized, the agent chooses an action x and pays $t(x \mid \hat{\boldsymbol{\theta}})$. In state u, he chooses x to maximize $u(x) - t(x \mid \hat{\boldsymbol{\theta}})$, whereas in state v, he chooses x to maximize $v(x) - t(x \mid \hat{\boldsymbol{\theta}})$.

We identify the direct mechanism with $t(x \mid \hat{\boldsymbol{\theta}})$, and we say that it implements the FB for a distribution of priors F if given this distribution, the game Γ has a PBNE such that for every profile of priors $\boldsymbol{\theta}$, expression (1) is equal to the FB surplus.

Proposition 1 *There exists a distribution F for which the FB is implementable, if and only if both $A - C > 0$ and $B - D > 0$. Moreover, as the ratio $\frac{A-C}{B-D}$ becomes closer to one, the set of distributions for which the FB is implementable expands. When $A - C = B - D$, the FB is implementable for every distribution F.*

Proof. Consider the class of mechanisms $t(x \mid \hat{\boldsymbol{\theta}})$, which satisfy condition (ii) in Remark 1 for every profile of reports $\hat{\boldsymbol{\theta}}$. For every such mechanism, it is optimal for the agent to play $x^u = a$ and $x^v = b$, regardless of the history. Our task is to show that there exists a mechanism in this class, such that if both players choose to participate in the mechanism and report their prior truthfully, we have a PBNE in the two-period game.

Apply the following affine transformations to the players' vNM utilities: *add* $\theta_1(B-D)$ to the speculator's vNM numbers, and *subtract* $(1-\theta_2)(B-C)$ from the agent's vNM numbers. Let $m_i(\hat{\boldsymbol{\theta}})$ denote player i's interim (transformed) expected payoff at the end of period 1, given that both players agreed to participate and that their report profile is $\hat{\boldsymbol{\theta}}$. Then:

$$\begin{aligned} m_1(\hat{\boldsymbol{\theta}}) &= \theta_1 \cdot t(a, \hat{\boldsymbol{\theta}}) + (1-\theta_1) \cdot t(b, \hat{\boldsymbol{\theta}}) + \theta_1(B-D) \\ m_2(\hat{\boldsymbol{\theta}}) &= \theta_2 \cdot [A - t(a, \hat{\boldsymbol{\theta}})] + (1-\theta_2) \cdot [B - t(b, \hat{\boldsymbol{\theta}})] - (1-\theta_2)(B-C) \end{aligned}$$

Because $t(x \mid \hat{\boldsymbol{\theta}})$ satisfies condition (ii) in Remark 1:

$$\begin{aligned} m_1(\hat{\boldsymbol{\theta}}) &= \begin{cases} \theta_1[(A-C)+(B-D)] + t(b, \hat{\boldsymbol{\theta}}) & if \quad \hat{\theta}_1 > \hat{\theta}_2 \\ t(b, \hat{\boldsymbol{\theta}}) & if \quad \hat{\theta}_1 < \hat{\theta}_2 \end{cases} \\ m_2(\hat{\boldsymbol{\theta}}) &= \begin{cases} -t(b, \hat{\boldsymbol{\theta}}) & if \quad \hat{\theta}_1 > \hat{\theta}_2 \\ \theta_2[(A-C)+(B-D)] - t(b, \hat{\boldsymbol{\theta}}) & if \quad \hat{\theta}_1 < \hat{\theta}_2 \end{cases} \end{aligned}$$

We may therefore rewrite $m_i(\hat{\boldsymbol{\theta}})$ as follows:

$$m_i(\hat{\boldsymbol{\theta}}) = \theta_i \cdot q_i(\hat{\boldsymbol{\theta}}) + \tau_i(\hat{\boldsymbol{\theta}})$$

where

$$q_i(\hat{\boldsymbol{\theta}}) = \begin{cases} (A-C)+(B-D) & if \quad \hat{\theta}_i > \hat{\theta}_j \\ 0 & if \quad \hat{\theta}_i < \hat{\theta}_j \end{cases} \tag{2}$$

and

$$\tau_1(\hat{\boldsymbol{\theta}}) = -\tau_2(\hat{\boldsymbol{\theta}}) = t(b, \hat{\boldsymbol{\theta}}) \tag{3}$$

Let $\bar{m}_i$ denote player i's expected (transformed) payoff, when at least one of the players refuses to participate in the mechanism. Then:

$$\begin{aligned} \bar{m}_1 &= \theta_1(B-D) \\ \bar{m}_2 &= \theta_2(A-C) \end{aligned}$$

It follows that our mechanism-design problem can be restated as follows. A partnership of size $(A-C)+(B-D)$ is jointly owned by two players, where player 1's initial share is $\frac{B-D}{(A-C)+(B-D)}$ and player 2's initial share is $\frac{A-C}{(A-C)+(B-D)}$. The value that each player i attaches to the partnership is θ_i, and it is common knowledge that θ_1 and θ_2 are independently drawn from a continuous *cdf* F with support $[0, 1]$. The problem is to design a direct mechanism that implements efficient dissolution of the partnership. A direct mechanism is identified with pair $(q_i(\hat{\boldsymbol{\theta}}), \tau_i(\hat{\boldsymbol{\theta}}))_{i=1,2}$, where $q_i(\hat{\boldsymbol{\theta}})$ satisfies (2) (ownership rights are fully transferred to the highest-reporting player), and $\tau_i(\hat{\boldsymbol{\theta}})$ satisfies (3) (a balanced budget condition). The question is whether there exists a PBNE, in which both players choose to participate and report truthfully, such that efficient dissolution is implemented.

If $A - C = 0$ or $B - D = 0$, the Myerson-Satterthwaite impossibility theorem applies. That is, there exists no distribution F, for which efficient dissolution is PBNE-implementable. If $A - C > 0$ or $B - D > 0$, we can apply Proposition 3 in CGK (1987). This result establishes that any two-person partnership which is not initially owned by a single agent can be dissolved efficiently for some distribution F. Proposition 1 in CGK (1987) shows that as initial shares become more equalized, the family of distributions F for which efficient dissolution is implementable expands. When initial shares are equal, efficient dissolution is implementable for any F. ■

The main lesson from this result is that implementability of optimal bets diminishes as the agent's incentives to manipulate the bet's outcome become more uneven across states. To develop an intuition for this result, it is helpful to make the following analogy. Think of a lottery in which a ball is drawn from an urn containing balls marked with the letter u and balls marked with the letter v. Neither the agent nor the speculator know the composition of the balls in the urn, but each has his own beliefs regarding the likelihood of drawing each type of ball. The party that draws the ball from the urn receives from the other party a monetary award if and only if the ball drawn is labeled u (otherwise, no payments are made). However, the agent can pay an amount of $(A - C) + (B - D)$ and guarantee that the lottery's outcome will be in his favor - that is, if he draws the ball it will be labeled u and if the speculator draws the ball it will be labeled v. Therefore, the speculator has no incentive to participate in this lottery if the monetary award exceeds $(A - C) + (B - D)$.

Assume the monetary award is $(A - C) + (B - D)$, and consider the problem of deciding who will draw the ball, the agent or the speculator. Assume that if the two parties cannot agree, then there is a third party that lets the agent draw a ball, and pays him $A - C$ if he draws a u ball. Thus, even if the agent can guarantee a favorable outcome in this lottery by paying the above amount $(A - C) + (B - D)$, he has no incentive to do so. The speculator, however, does not have this outside option.

The value that each party assigns to the right to draw a ball depends on his subjective prior probability of winning. It is interim-efficient to assign this right to the party with the highest prior. In order for the parties to agree to this assignment, the party who surrenders it to the other party needs to be compensated by the latter. In negotiating the terms of this compensation, each of the parties has some bargaining power. The agent can credibly threaten that he would veto any compensation scheme that leaves him worse off than his outside option (receiving $A-C$ if he draws a u ball). The speculator, on the other hand, can argue that without him, the monetary award would be $A-C$ and not $(A-C)+(B-D)$. Hence, the value of his consent is equal to the value of a lottery that pays a monetary award of $B-D$.

It follows that the efficient assignment of the right to draw a ball is equivalent to the following problem. The agent and the speculator can jointly produce $(A-C)+(B-D)$ units of a perfectly divisible asset. On his own, the agent can produce only $A-C$ units of this asset, while the speculator can produce none. The jointly produced asset is worth $\theta_1[(A-C)+(B-D)]$ to the speculator and $\theta_2[(A-C)+(B-D)]$ to the agent, where the values of θ_1 and θ_2 are privately and independently drawn from the same distribution. The question is, can we design a mechanism that ensures the assignment of the jointly produced asset to the party that values it the most, subject to the constraint that both parties have an incentive to produce this asset?

This is precisely the question asked by CGK in the context of an efficient dissolution of a partnership, where one partner initially holds $\frac{A-C}{(A-C)+(B-D)}$ of the partnership and another partner holds the remaining share. Thus, if $A-C \gg B-D$, such that the agent's initial share in the jointly owned "ball drawing rights" is close to one - in other words, if the agent enters the negotiation over the bet as a "seller" of the right to draw a ball, while the speculator enters as a "buyer" - the same forces that underlie the Myerson-Satterthwaite theorem make it hard to allocate the ball-drawing right to the player with the highest prior on u. As the gap between $A-C$ and $B-D$ shrinks, each player enters the negotiation both as a seller and a buyer, and thus he has what is often referred to as "countervailing incentives" when reporting his prior.

When $A-C=B-D$, the FB can be implemented using a natural *indirect* mechanism. Suppose that in period 1, the two players play a first-price, sealed-bid auction in order to determine which of them gets a lottery ticket that entitles its owner to a prize of $2(A-C)$ if and only if the agent chooses a in period 2. Both the revenues from the auction and the cost of paying the prize are equally distributed between the two players. Let Γ denote the two-period game induced by the auction.

Proposition 2 Γ *implements the FB for any distribution F. (We omit the proof because an analogous result with the same kind of proof is proved in Section 3.)*

Let us summarize the lessons from our results. When two players bet on outcomes which are susceptible to manipulation by one of them, an interim efficient bet is essentially an assignment of the right to receive an amount of money in one state. The amount is determined by the cost of manipulation (specifically, it is equal to the sum of the costs of manipulation in the two states). Efficient allocation of the right depends only on the players' beliefs. However, implementability of the efficient allocation diminishes as the costs of manipulation become more uneven across states. When they are infinitely uneven, we get a Myerson-Satterthwaite impossibility result. When they are completely even, we get a CGK possibility result.

Note that the model studied in this section assumes that if the two players do not sign any bet, the agent receives the entire surplus of the bare decision problem. Alternatively, we could assume that although the agent takes the action in period 2, all the surplus is reaped by the speculator (which would fit some variations on the buyer-seller story) None of our results would change. Similarly, our results would not change if we assumed that the second-period surplus is equally divided between the two players in each state.

3 Betting on future market prices

The previous section demonstrated the basic insight that efficient speculation over a manipulable outcome is analogous to an efficient dissolution of a partnership. In this section, we apply this insight to a market setting, in which traders with non-common prior beliefs regarding the size of future demand bet on future prices. As in Section 2, we study a two-period model. In period 2, the following market game is played. There are m sellers and n buyers. Each seller $s = 1, ..., m$ is able to supply a single unit of an indivisible good at a cost of $c \geq 0$. Sellers derive no utility from consuming the good. Each buyer $b = m+1, ..., m+n$ is willing to pay 1 for a single unit, and derives no utility from consuming additional units. There is also an external demand for ν units at a price of 1. External demand behaves stochastically. There are two states of nature, l (no external demand) and h (high external demand), such that $v = 0$ in state l and $v = h$ (abusing notation) in state h. We assume that $h > m$.

The market agents trade according to the following simultaneous-move procedure, adapted from Dubey (1982). Every agent (buyers and sellers alike) submits a *buy* *order*, consisting of a bid price and a number of demanded units. In addition, every seller submits a *sell order*, namely an ask price for the unit he is able to produce. Both bid and ask prices must lie in $[0, 1]$. The market price is the *highest* market-clearing price, given the aggregate supply and demand curves induced by the agents' buy and sell orders. If there exists no market-clearing price, the outcome is "no trade". If a seller turns out to purchase a unit from himself, he does *not* incur the production cost c. If the numbers of demanded and supplied units are unequal at the market-clearing price, we apply an efficient rationing rule.

Agents have quasi-linear utilities. A buyer's payoff is $\min(1, q^d) - pq^d$ if he ends up buying q^d units at a price p. A seller's payoff is $(p - c) \cdot q^s - pq^d$ if he ends up selling q^s units and buying q^d units at a price p (q^d gets the values $0, 1, 2, ...$, and q^s gets the values $0, 1$). Denote agent j's payoff function by u_j.

The realization of ν is common knowledge in period 2. In period 1, however, agents have conflicting prior beliefs regarding the likelihood of each state. Let θ_j denote the prior probability that agent j assigns to state h. For a seller s, θ_s can be interpreted as his degree of "optimism" regarding future external demand. Conversely, for a buyer b, θ_b is his degree of "pessimism" in this regard. Denote $\boldsymbol{\theta} = (\theta_j)_{j=1,\ldots,m+n}$. It is common knowledge that every agent j independently draws his prior belief θ_j from a continuous *cdf* F with support $[0, 1]$.

A bet is a multilateral contract, which maps a set of verifiable contingencies to budget-balanced monetary transfers among the n buyers and m sellers. We assume that neither the state of nature nor the agents' actions are verifiable. The only contingencies that can be contracted upon are whether trade occurs in the second period and at what price. For every action profile a in the market game, let $x(a) \in [0, 1] \cup D$ represent the verifiable market outcome induced by a, where $x(a) = D$ if a induces no trade, and $x(a)$ is the market price if a induces trade. Thus, a bet is a profile of functions $\mathbf{t} = (t_1(\cdot), ..., t_{m+n}(\cdot))$, where $t_j : [0, 1] \cup D \to \mathbb{R}$; $t_j(x)$ is the monetary transfer *received* by agent j when the second-period outcome is x; and $\sum_{i=1}^{m+n} t_i(x) = 0$ for all $x \in [0, 1] \cup D$.

If agents sign a bet in period 1, their second-period payoff function is modified, such that agent j's payoff from an action profile a is $u_j(a) + t_j(x(a))$. Consider an agent j who signed the bet and expects the second period action profile in state $\omega \in h, l$ to be a^ω. Denote $\mathbf{a} = (a^h, a^l)$ and $\mathbf{x} = (x(a^h), x(a^l))$. Given the agent's first-period prior belief, his expected utility is:

$$U_j(\mathbf{a}, \mathbf{t}) \equiv \theta_j[u_j(a^h) + t_j(x(a^h))] + (1 - \theta_j)[u_j(a^l) + t_j(x(a^l))]$$

We conclude the description of the model with a few comments:

- Second-period trade takes place once and for all. If an agent purchased a number of units which he cannot consume, he cannot resell them.

- Short-selling is ruled out: a seller cannot offer more than one unit and a buyer cannot offer any unit. Consequently, there is an asymmetry in the agents' ability to influence market outcomes. If the market price is $p < 1$, then every agent can unilaterally induce a higher price $p' \in (p, 1]$, by demanding a sufficiently large quantity at p'. In contrast, downward price manipulation is typically impossible, because a comparable "dumping" strategy is unavailable. The implication is that the only relevant second-period constraint facing agents as they sign bets will be to prevent unilateral upward manipulation of market prices.

- Observe that the assumptions imposed on the stochastic behavior of external demand simplify the SPIC constraints. Because the size of external demand in state h is higher than m, trade must take place in state h at a price of 1, regardless of the agents' actions. When agents contemplate signing a bet in period 1, they all agree that the verifiable market outcome in state h will be $x^h = 1$.

3.1 Interim-efficient bets

We now explore the limits that the agents' ability to manipulate market prices imposes on potential gains from speculative bets. Consider the following constrained optimization problem. For every profile of priors $\boldsymbol{\theta}$, choose a bet $\mathbf{t}(\boldsymbol{\theta})$ and a state-contingent action profile $\mathbf{a}(\boldsymbol{\theta})$ so as to maximize

$$\sum_j U_j[\mathbf{a}(\boldsymbol{\theta}), \mathbf{t}(\boldsymbol{\theta})] \tag{4}$$

subject to constraint that for every state $\omega \in \{h, l\}$, the outcome $a^\omega(\boldsymbol{\theta})$ is a Nash equilibrium in the modified market game in which agent j's payoff function is $u_j(a^\omega) + t_j(x(a^\omega))$. We call this SPIC, drawing on the terminology of Section 2. In order to be sustainable, a bet must satisfy the SPIC constraints - that is, it must provide the agents with incentives not to manipulate the market price.

A solution $(\mathbf{a}^*(\boldsymbol{\theta}), \mathbf{t}^*(\boldsymbol{\theta}))$ to the constrained optimization problem is *interim Pareto efficient*. In other words, for any pair $(\mathbf{a}, \mathbf{t})$ which is not a solution, the agents can find a bet $\mathbf{t}'$ and a state-contingent action profile $\mathbf{a}'$, such that every agent will prefer $(\mathbf{a}', \mathbf{t}')$ to $(\mathbf{a}, \mathbf{t})$, given his prior. We refer to the optimal value of (4) as the "FB surplus". Occasionally, we refer to set of solutions $(\mathbf{a}^*(\boldsymbol{\theta}), \mathbf{t}^*(\boldsymbol{\theta}))$ as "the FB" and to $\mathbf{t}^*(\boldsymbol{\theta})$ as an "optimal bet".

The following pair of examples illustrates how the SPIC constraints affect the sustainability of bets. In both examples, $m = n = 1$. The buyer is denoted b and the seller is denoted s. Our first example describes a bet which cannot be sustained, once SPIC constraints are taken into account. Suppose that b and s sign a bet requiring s (b) to pay A to his opponent if trade occurs (does not occur) in period 2. Thus, $t_s(D) = -t_b(D) = A$, and $t_s(x) = -t_b(x) = -A$ for every $x \in [0, 1]$. As we observed at the end of Section 2, occurrence of trade in state h is assured, regardless of the players' actions. Suppose that there were an action profile a^l such that $x(a^l) = D$. Then, the agents' first-period interim expected utilities would be:

$$
\begin{aligned}
U_s(\mathbf{a}, \mathbf{t}) &\equiv \theta_s \cdot [1 - c - A] + (1 - \theta_s) \cdot A \\
U_b(\mathbf{a}, \mathbf{t}) &\equiv \theta_b \cdot [1 - 1 + A] - (1 - \theta_b) \cdot A
\end{aligned}
$$

However, the buyer can impose trade in state l by demanding one unit at $p = 1$. Both before and after this deviation, his bare-game payoff is zero, but the deviation tilts the outcome of the bet in his favor. Therefore, as long as $A > 0$, there is no action profile that satisfy the SPIC constraints.

Now suppose that b and s sign an alternative bet requiring s to pay $p - c$ if there is trade at a price of $p > c$, and zero if there is no trade, or if there is trade at a price of $p \leq c$. This contract resembles a call option which is settled in cash, giving the buyer the right to purchase a unit of the good for a price of c in period 2. In state h trade occurs at $p = 1$, regardless of the agents' actions. Suppose that in state l, s (b) offers (demands) one unit at $p = c$. Let us show that this action profile constitutes a Nash equilibrium in the market game modified by the bet. The seller can manipulate the outcome of the bet only by raising the ask price to $p > c$ and demand at least one unit at the same price. However, his bare-game payoff will remain zero and in addition he will have to pay $p - c$ to the buyer. The buyer can manipulate the outcome by raising his bid price to p. The increase in the side payment that the buyer receives as a result of this deviation is exactly offset by the decrease in his bare-game payoff. Therefore, none of the agents wish to manipulate the bet's outcome. It follows that the bet and the constructed action profile satisfy the SPIC constraints.

Let us turn to the characterization of interim efficiency. First, we show that it does not compromise ex-post efficiency.

Proposition 3 *Let $(\mathbf{a}^*(\boldsymbol{\theta}), \mathbf{t}^*(\boldsymbol{\theta}))$ be interim efficient. Then, both a^h and a^l are ex-post efficient in the bare game (as well as in the modified game).*

This result is not self-evident. In order to have a non-trivial bet, there market outcome must vary across states. But when $n > m$, the competitive market price in the bare game is 1 in both states. Therefore, the price in state l is forced to be non-competitive. Therefore, the optimal bet must assume the role otherwise played by prices, namely providing the incentives needed for an efficient allocation. Our result shows that indeed this can be done.

Ex-post efficiency has the following implication. Let α_s^l and α_b^l denote the probabilities that a seller and a buyer trade in state l. Then, $\alpha_s^l = \min(1, \frac{n}{m})$ and $\alpha_b^l = \min(1, \frac{m}{n})$. Denote $i^* = \min_i \theta_i$ Given that F is continuous, we shall be able to ignore the case in which several agents share the same prior.

Proposition 4 *The FB satisfies the following properties.*

$(FB1)$ *If $m > n$, then $p^l = c$. If $m \leq n$, p^l can be any price in $[c, 1)$.*

$(FB2)$ *The FB surplus is equal to*

$$(1-c)\cdot\min\{m,n\} + (m-1)\cdot\sum_b(\theta_j - \min_i \theta_i) + [m-1+\alpha_s^l(1-c)]\cdot\sum_s(\theta_s - \min_i \theta_i) \quad (5)$$

$(FB3)$ *An optimal bet $\mathbf{t}^*(\boldsymbol{\theta})$ must satisfy:*

$$\begin{aligned} t_{s\neq i^*}(1) - t_{s\neq i^*}(p^l) &= \alpha_s^l(p^l - c) + m - 1 \\ t_{b\neq i^*}(1) - t_{b\neq i^*}(p^l) &= \alpha_b^l(1 - p^l) + m - 1 \end{aligned}$$

$(FB4)$ *The FB surplus can be achieved with the following class of bets:*

$$
\begin{aligned}
t_{s\neq i^*}(x) &= \begin{cases} T_s & for \quad p \leq p_* \ or \ D \\ T_s + \alpha_s^l(p^l - c) + pm - p & for \quad p_* < p \leq 1 \end{cases} \\
t_{b\neq i^*}(x) &= \begin{cases} T_b & for \quad p \leq p_* \ or \ D \\ T_b + \alpha_b^l(1 - p^l) + pm - 1 & for \quad p_* < p \leq 1 \end{cases} \\
t_{i^*}(x) &= -\sum_{j\neq i^*} t_j(x) \ for \ all \ x
\end{aligned}
$$

where T_i can be any agent-specific constant.

The FB constrains p^l to be competitive when there are more sellers than buyers. Otherwise, p^l can take any value in $[c, 1)$. Recall that $p^h = 1$ by our assumptions on external demand. Observe that p^l and p^h are independent of the profile of priors.

Turning to the FB surplus, the first term represents the bare-game surplus resulting from ex-post efficient trade. The second and third terms represent speculative gains. Under an optimal bet, the agent i^* with the lowest prior on h essentially bets on a low price (p^l) against each of his opponents j. The stakes of such a "bilateral bet" are constrained by the opponent's cost of manipulating the market price from p^l to p^h. The fact that only *upward* price manipulation is relevant in this regard originates from the assumption of no short-selling (see our discussion in Section 2).

For a buyer, manipulating the price upward in state l involves buying $m - 1$ unconsumed units at a price of $p = 1$. Note that in calculating this cost, we leave out the change in the price of the unit the consumer does consume. The intuition is that the changed price for the consumed unit is part of the transfer associated with the outcome $x = p^h$, rather than a cost involved in unilaterally inducing this outcome. The seller's cost of upward price manipulation involves two components. First, he needs to purchase $m - 1$ units which he does not consume, at a price of $p = 1$. Second, he purchases one unit from himself, thereby losing the gain from selling it to a buyer.

The speculative gain from each of the "bilateral bets" between i^* and j is equal to the stakes of the bet, multiplied by the difference between the agents' priors. The structure of the bet is simple and independent of the profile of priors. As long as trade occurs at a price $p \leq p^l$ or does not occur at all, the transfer is fixed. If the market price exceeds p^l, agent i^* pays an additional amount which is linear in the price.

3.2 Implementation

We now turn to the problem of finding a mechanism that implements the FB, when the agents' priors are their private information, and it is common knowledge that each agent independently draws his prior from F. A mechanism is a game played in the first period, whose outcome is a bet that modifies the payoff function of the second-period market game. We say that the mechanism weakly implements the FB if the two-period game induced by the mechanism has a Perfect Bayesian Nash Equilibrium (PBNE) in which the sum of the agents' interim expected utilities - *at the end of the first period and before the second-period game begins* - equals the FB

We require the mechanism to satisfy a participation constraint. Every agent can veto the mechanism, in which case the agents play a Nash equilibrium of the bare market game in period 2. Therefore, the interim expected utility that any agent earns in the PBNE of the two-stage game induced by the mechanism cannot be lower than his interim expected utility in the Nash equilibrium of the bare game. Note that when $m = n$, there are multiple equilibria in the bare game. In this case, the participation constraint is non-standard, in the sense that the agents' reservation utility is determined in equilibrium, rather than being exogenous.

Let us begin with the case of $m \neq n$, in which the bare-game equilibrium is unique. We consider implementation via a direct mechanism. This means that the agents play a two-period game, denoted Γ^*. In the first period, every agent submits a report $\hat{\theta}_j \in [0,1]$ or chooses to veto the mechanism. If all agents choose to participate, their profile of reports $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, ..., \hat{\theta}_{m+n})$ is assigned a profile of transfer functions $\mathbf{t}(x; \hat{\boldsymbol{\theta}}) = (t_1(x \mid \hat{\boldsymbol{\theta}}), ..., t_{m+n}(x \mid \hat{\boldsymbol{\theta}}))$. In period 2, the state of Nature is realized and the agents play the market game whose payoffs are modified by $\mathbf{t}(x \mid \hat{\boldsymbol{\theta}})$. We identify the direct mechanism with $\mathbf{t}(x \mid \hat{\boldsymbol{\theta}})$, and we say that it implements the FB for a distribution of prior beliefs F if given this distribution, the game Γ^* has a PBNE such that for every profile of priors $\boldsymbol{\theta}$, the sum of the agents' interim expected utility at the end of period 1 is equal to the FB surplus.

Proposition 5 *Let $m \neq n$. If $m = 1$, The FB is not implementable for any F. If $m > 1$, there exists a distribution F for which the FB is implementable.*

As in the previous section, the technical basis for this result is a formal analogy to the partnership dissolution model due to CGK. Implementability of efficient dissolution of the partnership diminishes as the initial ownership structure becomes more asymmetric. The analogous asymmetry in the model of Section 2 concerns the agent's cost of manipulating the bet's outcome in each state. In the current market model, price manipulation is relevant only in state l. However, buyers and sellers face different manipulation costs, as we saw in the discussion of expression (5).

The buyer-seller asymmetry plays an important role in Proposition 5. When $m = 1$, buyers face zero manipulation costs because they do not need to buy unconsumed units in order to drive prices up. In contrast, the seller has positive manipulation costs, because in order to drive the price up he needs to buy the good from himself and forego the income from selling to it a buyer. Therefore, we have an extreme buyer-seller asymmetry that results in a Myerson-Satterthwaite impossibility result. When $m > 1$, the asymmetry becomes less extreme, and we can apply the possibility result due to CGK.

Let us turn to comparative statics.

Proposition 6 *Fix F, and suppose that the FB is implementable for some m, n, c, $m \neq n$. Then:*
(i) The FB is also implementable for $c' \in (c, 1)$
(iii) The FB is also implementable for m', n' satisfying $m' > m$ and $n' \neq m'$.

As c tends to 1, manipulation costs become similar for buyers and sellers. Similarly, when the number of sellers becomes larger, the difference between buyers' and sellers' valuation of a single unit becomes negligible relative to the number of units that need to be purchased in order to drive the price up. Therefore, these changes in the market model's fundamentals facilitate implementability of the optimal bet.

The case of $m = n$ turns out to be special, because of the multiplicity of equilibria in the bare game. Consider the following indirect mechanism. In period 1, every agent chooses whether to exercise a veto option. If at least one agent exercises his veto option, the agents play the bare market game in the second period. If none of the agents exercise the veto option, the agents play a first-price, sealed-bid auction to determine which of them will receive a lottery ticket. The ticket entitles its owner to a prize of $Z = 2n \cdot (n - \frac{1+c}{2})$ if and only if trade occurs at a price $p < 1$. After the ticket is allocated to the winner, the agents play the market game in the second period. Both the revenues from the winner's bid and the cost of paying the prize are distributed equally among all agents.

The betting auction modifies second-period payoffs as follows (the bids are sunk at that stage, and therefore we can ignore them). If the market price is below 1, then in addition to the bare-game payoff, the auction winner receives a net payment of $(2n-1)\cdot(n-\frac{1+c}{2})$. The other agents' net payoff is their bare-game payoff minus $n-\frac{1+c}{2}$. If the market price is equal to 1, or if there is no trade, then the agents' net payoff is equal to their bare-game payoff. Let Γ denote the two-stage game induced by the betting auction.

Proposition 7 *Let $m=n$. Then, Γ implements the FB for all distributions F. Moreover, in the PBNE that implements the FB, $p^h=1$ and $p^l=\frac{1+c}{2}$ in period 2, regardless of the history.*

We have already commented on the two types of asymmetry between buyers and sellers. On one hand, manipulation-cost asymmetry means that driving the price from p^l up to $p^h=1$ is more costly for sellers than for buyers. When $m=n$, the cost differential is $1-c$. On the other hand, if $p^h>p^l$ in the bare-game equilibrium that is played whenever at least one agent exercises his veto option, this introduces an asymmetry in the opposite direction. However, if we select a bare-game equilibrium such that $p^h-p^l=\frac{1-c}{2}$, the two effects offset each other. Because any price in $[c,1]$ can be sustained in equilibrium in the bare game, we can set $p^l=\frac{1+c}{2}$, such that this condition is satisfied.

Note that the PBNE that implements the FB has the property that market prices are history-independent. In other words, the bets induced by the mechanism are "purely speculative", in the sense that they do not affect the outcome in the second-period market.

The upper bound on bid and ask prices

In the market model developed above, bid and ask prices are bounded in $[0,1]$. What is the economic justification for this assumption? Recall that the agents' valuations are common knowledge in the model. Therefore, it is also common knowledge that if an agent submits a bid price above 1, he must be exploiting his market power to tilt the outcome of a previously signed bet. An external regulatory agency may respond to such a transparent attempt to manipulate the price by shutting down the market, or by punishing the manipulator. This type of regulatory intervention is sometimes observed in sporting events and other contests.

Suppose that we relax this assumption, and allow agents to submit any non-negative bid and ask price. When $m=n=1$, this perturbation does not alter our analysis. The reason is that every agent can unilaterally impose no trade whenever the market price is strictly between 0 and 1. The SPIC constraints that follow are sufficiently strong to render the bounds on bid and ask prices irrelevant.

When there are more than two agents, relaxing the bound on prices implies that optimal bets do not exist, because the agents can sustain bets with arbitrarily high stakes. The trick is to set p^l between c and 1, and let p^h be arbitrarily high (note that external demand therefore becomes irrelevant). In this way, the cost of manipulating the price from p^l to p^h is also arbitrarily large, which allows agents to raise the stakes of their bet without limit.

Actually constructing a second-period Nash equilibrium that will sustain an arbitrarily high p^h is not trivial. The reason is that such a price exceeds the buyers' willingness to pay. If buyers could refrain from purchasing the good at $p^h > 1$ without affecting the market price, they would opt to do so. The equilibrium construction takes this into account: only one buyer b^* purchases the good in h. The bet is designed such that when $p = p^h$, some other market agent i gives b^* a transfer that compensates him for purchasing the good at $p^h > 1$. The reason i is willing to incur this cost is that he bets with a third agent that the second-period price will be p^h, and the speculative gains in this bilateral bet are sufficient to cover the compensatory transfer to b^*.

To conclude, while the bounds on market prices are irrelevant in the $m = n = 1$ case, they are crucial for our results when there are more than two agents.

The domain of bets

Our model assumes that bets are a function of whether trade occurs, and at which price. What would happen if bets could be a function of the agents' entire action profile? Once again, there is a difference between the case of $m = n = 1$ and the case of more than two agents. In the former case, it can be shown that our analysis continues to hold. The intuition is as follows. Let (a_s^h, a_b^h) and (a_s^l, a_b^l) denote the equilibrium action profiles in states h and l. Among the deviations that need to be prevented in order to sustain a bet, are: (i) the seller's deviation in h and the buyer's deviation in l into the profile (a_s^l, a_b^h), and (ii) the seller's deviation in l and the buyer's deviation in h into the profile (a_s^h, a_b^l). These constraints alone place an upper bound on the sum of the agents' interim expected utilities, which cannot exceed the FB surplus. Thus, the agents cannot do better than signing a bet that conditions only on whether trade occurs, and at what price.

In contrast, when there are more than two agents, allowing the agents to condition bets on action profiles gives rise to bets with arbitrarily large stakes. The reason is that agents i and j can bet on the action that agent k will take. Because i and j take k's action as given in Nash equilibrium, they cannot manipulate it and therefore nothing prevents them from raising the stakes of their bet without limit. We only need to ensure that $a_k^h \neq a_k^l$, in order for the bet to be possible. But this can be achieved because the speculative gains earned by i and j are sufficient to provide k with an incentive to take different actions in the two states.

To conclude, when $m = n = 1$, the assumption that bets are a function of a coarse partition of the set of possible outcomes is not necessary for our results. When there are more than two agents, this assumption is crucial.

An alternative model of demand uncertainty
Our market model assumes that demand fluctuation takes the form of "noise traders", whose presence in the market is uncertain. Alternatively, we can assume that there are no noise traders, but the buyers' value of the good is (abusing notation) h in state h and l in state l, with $h > l$.

We have analyzed this alternative model in the special case of $m = n = 1$. The FB is ex-post efficient. When $h > l > c$, the second-period equilibrium has the property that $p^h > p^l$ if $\theta_s > \theta_b$ and $p^h < p^l$ if $\theta_s < \theta_b$ - that is, the assignment of prices to states is sensitive to the priors. The FB is implementable (via a variant on the above "betting auction" mechanism for any distribution F. The reason is that as in the "noise traders" version of the model, when $m = n = 1$ the "bare" market game has a continuum of equilibria.

4 Discussion

High-order uncertainty. In our models, the distribution from which the agents independently draw their priors is common knowledge. Thus, we assume away high-order uncertainty regarding the agents' prior beliefs. One may argue that if agents cannot agree on their first-order beliefs, it is hard to imagine that they can agree on high-order beliefs. In addition, Yildiz (2004) argues that the coexistence of Nash equilibrium and non-common priors over states of nature is methodologically problematic, because in reality it is usually more difficult to predict strategic behavior than the behavior of "nature". Our motivation for this assumption is primarily technical: we wish to parallel the benchmark mechanism-design models in the literature. In some contexts we find the assumption reasonable. For instance, in the context of the "central bank" story of Section 2, high-order certainty approximates the fact that opinions regarding the future rate of inflation are often surveyed and made public.

Non-common priors vs. state-dependent utility. Our model is formally indistinguishable from a model in which every agent j assigns probability $\frac{1}{2}$ to each state, and his utility function is multiplied by a state-dependent constant (θ_j in one state and $1 - \theta_j$ in the other state). The motivation for signing side contracts under this re-interpretation is risk sharing rather than speculative trade. Note that this re-interpretation requires us to assume that the utility from money is state-dependent, whereas the trade-off between money and consumption is state-independent. It is hard to imagine other motivations than non-common priors for such preferences. Of course, as we have seen, the formal analogy between our mechanism-design approach to speculative trade and models of preference-based trade proved analytically useful.

Related literature. A distinctive feature of our model is the focus on bets made between parties who can manipulate the bet's outcome. Bets are essentially side transfers that modify the payoffs of the second-period game. A similar insight was used by Allaz and Vila (1993) to derive a rationale for forward markets, in an environment without uncertainty. They show that producers may wish to use forward contracts in order to improve their situation in a future, imperfectly competitive spot market. In their model, producers first trade in forward contracts, and then play a Cournot game in which their payoff functions are modified by the positions they took in the forward market. The authors assume that the forward market is perfectly competitive and includes traders who are not active in the spot market. Dong and Liu (2005) study an imperfectly competitive forward market in which all traders are also actors in a competitive spot market.

Wilson (1968) investigates the problem faced by a group of agents who need to make a collective decision that generates a surplus whose value depends on an uncertain state of nature. The question is, how should this surplus be divided among the agents in order to ensure Pareto optimality of the collective decision? Wilson allows for non-common priors. Therefore, efficient sharing rules may involve side bets on the value of future surplus. The outcome of these bets can be manipulated by the agents, because the surplus depends on the collective decision that is made. Wilson (1968) provides a necessary and sufficient condition for Pareto optimality of a sharing rule, and gives examples of such rules in specific environments.

Finally, the present paper follows up Eliaz and Spiegler (2004,2005). These papers analyze the problem of designing a profit-maximizing menu of contracts for a monopolist who faces consumers who differ in their ability to forecast their future tastes. In Eliaz and Spiegler (2004), the agent's preferences are dynamically inconsistent, and agent types differ in the prior probability they assign to the possibility that their tastes will not change (interpreted as their degree of naivete). Eliaz and Spiegler (2005) analyze a similar problem with dynamically consistent preferences. Both papers study environments in which non-common priors turn out to be necessary for price discrimination.

References

[1] Allaz, Blaise and Jean-Luc Vila (1993): "Cournot Competition, Forward Markets and Efficiency," *Journal of Economic Theory*, 59, 1-16.

[2] Cramton, Peter, Robert Gibbons and Paul Klemperer (1987): "Dissolving a Partnership Efficiently," *Econometrica,* 55, 615-632.

[3] Chung, Kim-Sau and Jeffrey Ely (2005): "Foundations of Dominant Strategy Mechanisms", Working Paper, Northwestern University.

[4] Dong, Lingxiu and Hong Liu (2005): "Equilibrium Forward Contracts on Nonstorable Commodities in the Presence of Market Power," Working Paper, Olin School of Business, Washington University in St. Louis.

[5] Dubey, Pradeep (1982): "Price-Quantity Strategic Market Games," *Econometrica,* 50, 111-126.

[6] Eliaz, Kfir and Ran Spiegler (2004): "Contracting with Diversely Naive Agents," Working Paper, New York University and Tel Aviv University.

[7] Eliaz, Kfir and Ran Spiegler (2005): "Speculative Contracts," Working Paper, New York University and Tel Aviv University.

[8] Milgrom, Paul and Nancy Stokey (1982): "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17-27.

[9] Morris, Stephen (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica,* 62, 1327-1347.

[10] Myerson, Roger and Mark Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265-281.

[11] Wilson, Robert (1968): "The Theory of Syndicates," *Econometrica,* 36, 119-132.

[12] Yildiz, Muhamet (2004): "Wishful Thinking in Strategic Environments," Mimeo, MIT.

Appendix: Proofs

Proof of Proposition 3

Lemma 1 *The outcome in state h ex-post efficient.*

Proof. In state h, the market price is $p^h = 1$, regardless of the agents' actions. Therefore, in equilibrium they will act as price takers: each seller will offer one unit and demand zero units, while each buyer is indifferent between demanding one unit and demanding zero units. Therefore, the outcome is ex-post efficient. ■

Lemma 2 *Trade occurs in state l.*

Proof. Assume the contrary. Each agent can manipulate the outcome and impose trade at $p = 1$, by demanding a single unit at a price of 1. Moreover, each seller can impose this at no cost by simultaneously submitting a bid of 1 and and an ask of 0, in which case he would buy the good from himself (it must be the case that all the other sellers quote a strictly positive ask price - otherwise, trade would occur). It follows that the SPIC constraints in the no-trade state must include the following inequalities:

$$\begin{aligned} t_s(D) &\geq 1 - 1 + t_s(1) \\ t_b(D) &\geq 1 - 1 + t_b(1) \end{aligned}$$

By budget balancedness, $\Sigma_i t_i(D) = \Sigma_i t_i(1) = 0$. Hence, $t_i(D) = t_i(1)$ for all i, such that total surplus is equal to the bare-game surplus, given the agents' behavior. But since the bare-game outcome is ex-post inefficient in state l, it obviously does not maximize total surplus. ■

Lemma 3 *If $p^l > c$, then every seller offers his unit. If $p^l < 1$, every buyer demands at least one unit.*

Proof. If $p^l > c$, an inactive seller s can be active by offering his unit at $p = p^l$. This will not change the market price, and will increase his payoff by $p - c$ with positive probability. If $p^l < 1$, an inactive buyer b can be active by demanding one unit at $p = p^l$. This will not change the market price, and his payoff will increase by $1 - p$ with positive probability. ■

Lemma 4 $p^l \geq c$.

Proof. Assume the contrary. Then, there must be exactly one seller s^* who actually sells his unit - otherwise, any individual seller could deviate by asking a price above c, and he would avoid the unprofitable exchange without affecting the market price. In addition, there is no buyer who bids above p - otherwise, the market price would exceed p, because aggregate supply consists of exactly one unit. It follows that every agent can unilaterally impose any price $p' > p$, by demanding one unit at p'. In particular, every seller can impose p' by demanding one unit at p' while submitting an ask price below p.

Set $p' = c$ and let us write down the SPIC constraints induced by the above analysis:

$$\begin{aligned} p - c + t_{s^*}(p) &\geq t_{s^*}(c) \\ t_s(p) &\geq t_s(c) \text{ for every } s \neq s^* \\ \alpha_b \cdot (1-p) + t_b(p) &\geq 1 - c + t_b(c) \text{ for every } b \end{aligned}$$

where α_b is the probability that buyer b will consume one unit. Note that $\Sigma_b \alpha_b \leq 1$ (possibly with inequality, because some buyers may demand more units than others). Summing these inequalities yields $p - c + 1 - p \geq n \cdot (1-c)$, a contradiction. ■

If $p^l = 1$, then the market outcome is the same in both states. Therefore, total surplus is equal to the bare-game surplus. As we show in Proposition 4, the agents can attain a strictly higher surplus than the maximal bare-game surplus. Therefore, $p^l = 1$ cannot be induced by an optimal bet.

If $p^l \in (c, 1)$, then by Lemma 3, each seller offers his unit, and each buyer demands at least one unit. By efficient rationing, the outcome is ex-post efficient in state l.

If $p^l = c$, then by Lemma 3, each buyer demands at least one unit. The outcome can be inefficient only if some sellers refrain from offering their unit. Then, we can modify the action profile such that each seller offers his unit. This would not add new SPIC constraints, and it can only relax existing SPIC constraints, thereby increasing total surplus. ■

Proof of Proposition 4

By Proposition 3, the FB is ex-post efficient. We already know that $p^h = 1$ and $p^l \in [c, 1]$. In state h, the agents cannot affect the market outcome, and they behave as price takers. Our focus will thus be on the action profile in state l. Ignore for the moment SPIC constraints that concern manipulation of the market price downward. The following SPIC constraints - concerning the agents' ability to manipulate the price upward - must hold:

$$\alpha_s^l(p^l - c) + t_s(p^l) \geq p - mp + t_s(p) \quad (6)$$

$$\alpha_b^l(1 - p^l) + t_b(p^l) \geq 1 - mp + t_b(p) \quad (7)$$

for all $p > p^l$.

If $\min\{m, n\} = 1$, then there is at least one agent who can unilaterally impose no-trade in state l, hence there are additional SPIC constraints that prevent such a deviation. We can minimize these constraints by having all agents quote the same price. Hence, if $m = 1$ and $n > 1$, or if $n = 1$ and $m > 1$, only a single agent - either the single buyer or the single seller - can impose no trade. To prevent him from doing so, we can impose an infinite fine on him whenever there is no trade. That is, if $m = 1$ and $n > 1$ we set $t_s(D) = -\infty$, and if $n = 1$ and $m > 1$, we set $t_b(D) = -\infty$.

If $m = n = 1$, then each agent can unilaterally impose no trade in every state. This means that we need to satisfy additional SPIC constraints:

$$p^l - c + t_s(p^l) \geq t_s(D) \quad (8)$$

$$1 - p^l + t_b(p^l) \geq t_b(D) \quad (9)$$

Note that when $m = n = 1$ the SPIC constraints (6) and (7) with respect to $p = 1$ become

$$p^l - c + t_s(p^l) \geq t_s(1) \quad (10)$$

$$1 - p^l + t_b(p^l) \geq t_b(1) \quad (11)$$

Hence, by setting $t_s(D) = t_s(1)$ and $t_b(D) = t_b(1)$, we make the constraints (8) and (9) equivalent to (10) and (11). It follows that the additional constraints required to prevent no-trade when $\min\{m, n\} = 1$, can be satisfied without imposing further restrictions on $t_j(p)$, beyond those implied by (6) and (7).

For each $p > p^l$ define,

$$z_s(p) \equiv \alpha_s^l p + t_s(p) - \alpha_s^l p^l - t_s(p^l)$$

$$z_b(p) \equiv -\alpha_b^l p + t_b(p) + \alpha_b^l p^l - t_b(p^l)$$

The SPIC constraints can then be written more compactly as follows:

$$z_s(p) \leq (m + \alpha_s^l - 1)p - \alpha_s^l c \quad (12)$$

$$z_b(p) \leq (m - \alpha_b^l)p + \alpha_b^l - 1 \quad (13)$$

By budget-balancedness,

$$\begin{aligned} -z_s(p) &\leq [(m-1)(m+\alpha_s^l-1)+n(m-\alpha_b^l)]p - \alpha_s^l(m-1)c - n(1-\alpha_b^l) \\ -z_b(p) &\leq [m(m+\alpha_s^l-1)+(n-1)(m-\alpha_b^l)]p - \alpha_s^l mc - (n-1)(1-\alpha_b^l) \end{aligned}$$

Hence, the SPIC constraints imply that for all $p > p^l$,

$$\begin{aligned} z_s(p) &\geq -[(m-1)(m+\alpha_s^l-1)+n(m-\alpha_b^l)]p + \alpha_s^l(m-1)c + n(1-\alpha_b^l) \\ z_b(p) &\geq -[m(m+\alpha_s^l-1)+(n-1)(m-\alpha_b^l)]p + \alpha_s^l mc + (n-1)(1-\alpha_b^l) \end{aligned}$$

Total expected surplus is given by

$$\begin{aligned} &\sum_s \{\pi_s[\alpha_s^l(p^h - c) + t_s(p^h)] + (1-\pi_s)[\alpha_s^l(p^l - c) + t_s(p^l)]\} \\ &+ \sum_b \{\pi_b[\alpha_b^l(1 - p^h) + t_b(p^h)] + (1-\pi_b)[\alpha_b^l(1 - p^l) + t_b(p^l)]\} \end{aligned}$$

which (because $p^h = 1$) may be written more compactly as,

$$\sum_j \pi_j z_j(1) + \min\{m, n\} \cdot (1 - c)$$

By budget balancedness, we may write the above expression for the total surplus as follows:

$$\sum_{j \neq i^*} (\pi_j - \pi_{i^*}) z_j(1) + \min\{m, n\} \cdot (1 - c)$$

By (12) and (13), the surplus is bounded from above by

$$\begin{aligned} &\sum_{s \neq i^*} (\pi_s - \pi_{i^*})[(m-1) + \alpha_s^l(1-c)] + \sum_{b \neq i^*} (\pi_b - \pi_{i^*})(m-1) + \min\{m,n\} \cdot (1-c) \\ = \; & (m-1) \sum_{j \neq i^*} (\pi_j - \pi_{i^*}) + \alpha_s^l(1-c) \sum_{s \neq i^*} (\pi_s - \pi_{i^*}) + \min\{m,n\} \cdot (1-c) \end{aligned}$$

Note that the upper bound is independent of the value of p^l. In order to attain the upper bound, the constraints (12) and (13) must be binding for $p = 1$. This yields the result $FB3$. It remains to show that there exist a price $p^l \in [c, 1]$ and transfer functions that attain the upper bound on surplus. For each seller $s \neq i^*$, let

$$t_s(p) = \begin{cases} t_s(p^l) & for \quad p \leq p^l \\ t_s(p^l) + \alpha_s(p^l - c) + p(m-1) & for \quad p^l < p \leq 1 \end{cases}$$

and for each buyer $b \neq i^*$, let

$$t_b(p) = \begin{cases} t_b(p^l) & for \quad p \leq p^l \\ t_b(p^l) + \alpha_b(1 - p^l) + mp - 1 & for \quad p^l < p \leq 1 \end{cases}$$

To maintain budget balancedness, set

$$t_{i^*}(p) = -\sum_{j \neq i^*} t_j(p)$$

for every price p. By construction, the SPIC constraints are satisfied and the upper bound on surplus is met, as long as $p^l < 1$.

It remains to examine the SPIC constraints that concern downward price manipulation. When $n \geq m$, the above-constructed transfer function necessarily satisfies these constraints, for every $p^l \in [c, 1)$. When $m > n$, we can set $p^l = c$, such that these constraints are satisfied. ■

Proof of Proposition 5

Define $i_*(\hat{\boldsymbol{\theta}})$ to be the lowest indexed agent among those agents with the lowest reported prior on h. Consider a direct mechanism $\mathbf{t}(x; \hat{\boldsymbol{\theta}})$ defined as the FB bet (given by ($FB4$) in Proposition 4) as if $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$. Note that the value of T_i is left undetermined, and so we shall allow it be a function of $\hat{\boldsymbol{\theta}}$. We ask if for some distributions F, there exists a direct mechanism such that in the two-stage game it induces, there is a PBNE which satisfies (IR) and (IC), and induces $\mathbf{t}(x; \hat{\boldsymbol{\theta}})$. Thus, the bets signed under different profiles of reported priors differ only in their T.

Let us first construct second-period continuation strategies, treating the cases $n > m$ and $n < m$ separately.

Case 1: $n > m$

If at least one agent refuses to participate in the first-period mechanism, each seller offers one unit and each buyer demands one unit at a price of 1. This state-contingent action profile constitutes a NE in the bare game, hence in the second-period subgame in question. Now suppose that all agents agreed to participate in the first-period mechanism. In state h, each seller offers his unit at any price, and buyers behave arbitrarily. In state l, each seller offers one unit and each buyer demands one unit at a price $p^l \in [c, 1)$. As demonstrated in Proposition 4, this action profile constitutes a NE in the game modified by $\mathbf{t}(x; \hat{\boldsymbol{\theta}})$. Moreover, it induces the FB surplus. Although by construction $\alpha_s^l = 1$ and $\alpha_b^l = \frac{m}{n}$ in this case, for the sake of transparency we will use the more agnostic notation α_s^l and α_b^l (observing that $\Sigma_b \alpha_b^l = \Sigma_s \alpha_s^l = m$).

Case 2: $n < m$

If at least one agent refuses to participate in the first-period mechanism, then in state h, each seller offers one unit at an arbitrary price and each buyer behaves arbitrarily. In state l, each seller offers one unit and each buyer demands one unit at a price of c. This state-contingent action profile constitutes a NE in the bare game, hence in the second-period subgame in question. Now suppose that all agents agreed to participate in the first-period mechanism. In state h, each seller offers his unit at any price, and buyers behave arbitrarily. In state l, each seller offers one unit and each buyer demands one unit at a price $p^l = c$. This action profile constitutes a NE in the game modified by $\mathbf{t}(x; \hat{\boldsymbol{\theta}})$. Moreover, it induces the FB surplus. Although by construction $\alpha_s^l = \frac{n}{m}$ and $\alpha_b^l = 1$ in this case, for the sake of transparency we will use the more agnostic notation α_s^l and α_b^l (observing that $\Sigma_b \alpha_b^l = \Sigma_s \alpha_s^l = n$).

Our objective is thus to examine whether there exist distributions F for which agreeing to participate in the mechanism and reporting one's true prior, together with the second-period continuation strategies described above, constitute a PBNE. Apply the following affine transformation to the agents' utilities:

$$v_s \equiv u_s + (1 - \theta_s)(m - c) - (1 - c) \tag{14}$$

$$v_b \equiv u_b + (1 - \theta_b)(m - 1) \tag{15}$$

Denote:

$$\begin{aligned}
q_s(\hat{\boldsymbol{\theta}}) &\equiv \alpha_s^l (p^l - c) + (m - 1) - [t_s(1; \hat{\boldsymbol{\theta}}) - t_s(p^l; \hat{\boldsymbol{\theta}})] \\
q_b(\hat{\boldsymbol{\theta}}) &\equiv \alpha_b^l (1 - p^l) + (m - 1) - [t_b(1; \hat{\boldsymbol{\theta}}) - t_b(p^l; \hat{\boldsymbol{\theta}})] \\
\tau_s(\hat{\boldsymbol{\theta}}) &= t_s^h(\hat{\boldsymbol{\theta}}) \\
\tau_b(\hat{\boldsymbol{\theta}}) &= t_b^h(\hat{\boldsymbol{\theta}})
\end{aligned}$$

In addition, let $V_j(\mathbf{x}^*, \mathbf{t}^*; \hat{\boldsymbol{\theta}})$ be the expectation of v_j, given$(\mathbf{x}^*, \mathbf{t}^*; \hat{\boldsymbol{\theta}})$. Then, we can write:

$$v_i(\hat{\boldsymbol{\theta}}) = (1-\theta_i)\cdot q_i(\hat{\boldsymbol{\theta}}) + \tau_i(\hat{\boldsymbol{\theta}})$$

When $n > m$, if at least one agent refuses to participate in the first-period mechanism, then the expectation of u_s is $1-c$, and the expectation of u_b is 0. Let $\bar{V}_j(\theta_j)$ denote the expectation of v_j in this case. Then:

$$\begin{aligned}\bar{V}_s(\theta_s) &= (1-\theta_s)(m-c)\\ \bar{V}_b(\theta_b) &= (1-\theta_b)(m-1)\end{aligned}$$

Denote

$$\begin{aligned}\bar{q}_s(\theta_s) &\equiv m-c\\ \bar{q}_b(\theta_b) &\equiv m-1\end{aligned}$$

Then, $\sum_j \bar{q}_j(\theta_j) = m(m-c)+n(m-1)$.

When $n < m$, if at least one agent refuses to participate in the first-period mechanism, then the expectation of u_s is $\theta_s(1-c)$, and the expectation of u_b is $(1-\theta_b)(1-c)$. Let $\bar{V}_j(\theta_j)$ denote the expectation of v_j in this case. Then:

$$\begin{aligned}\bar{V}_s(\theta_s) &= (1-\theta_s)(m-1)\\ \bar{V}_b(\theta_b) &= (1-\theta_b)(m-c)\end{aligned}$$

Denote

$$\begin{aligned}\bar{q}_s(\theta_s) &\equiv m-1\\ \bar{q}_b(\theta_b) &\equiv m-c\end{aligned}$$

Then, $\sum_j \bar{q}_j(\theta_j) = m(m-1)+n(m-c)$.

Letting $E_{-j}(\cdot)$ be the expectation operator with respect to θ_{-j}, and define

$$V(\hat{\theta}_j,\theta_j) \equiv E_{-j}[(1-\theta_i)\cdot q_j(\hat{\theta}_j,\theta_{-j}) + \tau_j(\hat{\theta}_j,\theta_{-j})]$$

Our mechanism design problem is thus reduced to the problem of finding a pair of profiles, $(q(\hat{\boldsymbol{\theta}}))_i$ and $(\tau(\hat{\boldsymbol{\theta}}))_i$, that satisfy the following properties:

(EFF) for every agent j, and for every profile of reports $\hat{\boldsymbol{\theta}}$,

$$q_j(\hat{\boldsymbol{\theta}}) = \begin{cases} m(m-c)+n(m-1) & if \quad j = i^*(\hat{\boldsymbol{\theta}}) \\ 0 & if \quad j \neq i^*(\hat{\boldsymbol{\theta}}) \end{cases}$$

if $n > m$, and

$$q_j(\hat{\boldsymbol{\theta}}) = \begin{cases} m(m-1)+n(m-c) & if \quad j = i^*(\hat{\boldsymbol{\theta}}) \\ 0 & if \quad j \neq i^*(\hat{\boldsymbol{\theta}}) \end{cases}$$

if $n < m$.

(IC) for every agent j and for every θ_j,

$$V(\theta_j, \theta_j) \geq V(\hat{\theta}_j, \theta_j)$$

for all possible reports $\hat{\theta}_j$

(IR) for every agent j and for every θ_j,

$$V(\theta_j, \theta_j) \geq \bar{V}_j(\theta_j)$$

But this is precisely the problem of efficiently dissolving a partnership of $m+n$ agents. When $n > m$, each of m partners owns a fraction

$$r_s(m,n,c) \equiv \frac{m-c}{m(m-c)+n(m-1)} \tag{16}$$

and each of n partners owns a fraction

$$r_b(m,n,c) \equiv \frac{m-1}{m(m-c)+n(m-1)} \tag{17}$$

When $n < m$, each of m partners owns a fraction

$$r_s(m,n,c) \equiv \frac{m-1}{m(m-1)+n(m-c)} \tag{18}$$

and each of n partners owns a fraction

$$r_b(m,n,c) \equiv \frac{m-c}{m(m-1)+n(m-c)} \tag{19}$$

In both cases, each partner j draws his value $1-\theta_j$ of the entire partnership independently from a common continuous distribution. Note that if $m = 1$, then this partnership is owned entirely by the single seller. By Proposition 2 of CGK, a one-owner partnership cannot be dissolved efficiently for *any* distribution F. But if $m > 1$, then by Proposition 3 of CGK, any partnership not owned by a single agent can be dissolved efficiently for some distribution F. ■

Proof of Proposition 6

By Proposition 5, the implementation problem is equivalent to the efficient dissolution of a partnership with initial shares given by (16) and (17) when $n > m$, and by (18) and (19) when $n < m$. A simple glance at the formulas shows that $r_s(m, n, c)/r_b(m, n, c)$ becomes closer to one when: *(i)* the number of sellers increases (as long as $m \neq n$); *(ii)* c increases and becomes closer to one. By Proposition 1 of CGK, if the original partnership can be efficiently dissolved for some F, then so are the modified partnerships. ■

Proof of Proposition 7

Regardless of the state and of the history, assume that in period 2, each seller offers one unit and each buyer demands one unit. In state h, their ask/bid prices can be chosen arbitrarily, and in state l, they all quote the same price $p^l = \frac{1+c}{2}$.

This action profile constitutes a Nash equilibrium in the second-period subgame. Consider first the case of $m = n > 1$. In state h, no agent can unilaterally alter the outcome. In state l, an agent can alter the outcome of the bet by demanding n units at a price of 1. In order to prevent him from doing so, the following conditions must hold for every seller s and every buyer b. If the seller won the betting auction, he must satisfy $\frac{1+c}{2} - c + Z \geq 1 - n$. If he lost, he must satisfy $\frac{1+c}{2} - c - \frac{Z}{2n} \geq 1 - n$. Similarly, if the buyer won the betting auction, he must satisfy $1 - \frac{1+c}{2} + Z \geq 1 - n$, and if he lost, he must satisfy $1 - \frac{1+c}{2} - \frac{Z}{2n} \geq 1 - n$. Note that the first-period bids are sunk, hence they are left out of these constraints. Given the value of Z, it is easy to check that all four constraints hold.

When $m = n = 1$, we need to consider the agents' ability to impose no trade in state l. If the seller won the betting auction, the additional SPIC constraints are:

$$\begin{aligned} 1 - \frac{1+c}{2} - Z &\geq 0 \\ \frac{1+c}{2} - c + Z &\geq 0 \end{aligned}$$

and if the buyer won the auction, the constraints are:

$$\begin{aligned} 1-\frac{1+c}{2}+Z &\geq 0 \\ \frac{1+c}{2}-c-Z &\geq 0 \end{aligned}$$

Because $Z=\frac{1-c}{2}$ when $n=1$, these additional constraints are satisfied.

Suppose all agents expect the above NE to be played in the second period. Suppose also that for all realizations of priors, the agent with the highest prior on l wins (because F is atomless, we can ignore ties). It is easy to check that according to the mechanism, $t_i(p^h)-t_i(p^l)$ is consistent with the specification of optimal bets given by Proposition 4. Therefore, the sum of the agents' expected payoffs at the end of the first-period auction is equal to the FB surplus.

It remains to show that if agents expect to play the second-period NE described above, then the first-period auction has a BNE with the following properties: (1) no agent chooses to exercise his veto option; (2) the winning agent assigns the highest prior to state l.

Let $\lambda_j(\theta_j)$ denote the probability that agent j wins the auction, given that his prior on state h is θ_j. The interim expected payoff of a seller s, given a profile of bidding strategies $(\beta_j)_j$, is given by

$$\lambda_j(\theta_j)U_{j=i^*}(\mathbf{x}^*,\mathbf{t}^*)+[1-\lambda_j(\theta_j)]\cdot E_{\theta_{-j}}[U_{j\neq i^*}(\mathbf{x}^*,\mathbf{t}^*)|\ j\neq i^*] \tag{20}$$

where:

$$\begin{aligned} U_{s\neq i^*}(\mathbf{x}^*,\mathbf{t}^*) &= (1-\theta_s)(\frac{1+c}{2}-c-\frac{Z}{2n})+\theta_s(1-c)+\frac{\beta_{i^*}}{2n} \\ U_{s=i^*}(\mathbf{x}^*,\mathbf{t}^*) &= (1-\theta_{i^*})[\frac{1+c}{2}-c+(1-\frac{1}{2n})Z]+\theta_{i^*}(1-c)-(1-\frac{1}{2n})\beta_{i^*} \\ U_{b\neq i^*}(\mathbf{x}^*,\mathbf{t}^*) &= (1-\theta_b)(1-\frac{1+c}{2}-\frac{Z}{2n})+\theta_s(1-1)+\frac{\beta_{i^*}}{2n} \\ U_{b=i^*}(\mathbf{x}^*,\mathbf{t}^*) &= (1-\theta_{i^*})[1-\frac{1+c}{2}+(1-\frac{1}{2n})Z]+\theta_{i^*}(1-1)-(1-\frac{1}{2n})\beta_{i^*} \end{aligned}$$

If at least one agent vetoes the auction, then the seller's interim expected payoff is

$$U_s(\mathbf{x}^*,\mathbf{t}^0)=\theta_s(1-c)+(1-\theta_s)(\frac{1+c}{2}-c)=(1-\theta_s)(\frac{c-1}{2})+(1-c)$$

while the buyer's interim expected payoff is

$$U_b(\mathbf{x}^*, \mathbf{t}^0) = (1-\theta_b)(1-\frac{1+c}{2})$$

Apply the following affine transformations to the agents' vNM utility functions:

$$\begin{aligned} v_s &= u_s + (1-\theta_s)(n-c) - (1-c) \\ v_b &= u_b + (1-\theta_b)(n-1) \end{aligned}$$

Applying this transformation to the expected payoff of agent j from his bidding strategy $\beta_j(\theta_j)$ yields

$$\lambda_j(\theta_j) \cdot [(1-\theta_j)Z - \beta_j] + [1-\lambda_j(\theta_j)]\frac{1}{2n}E_{\theta_{-j}}[\beta_{i^*}(\theta_{i^*})|\ j \neq i^*] \tag{21}$$

where $Z = 2n(n - \frac{1+c}{2})$. Applying the same transformation to each agent's "default" payoff yields

$$V_j(\mathbf{x}^*, \mathbf{t}^0) = (1-\theta_j)\frac{Z}{2n} \tag{22}$$

We wish to show that there exists a profile of monotonic bidding strategies $(\beta_j(\theta_j))_j$ that constitute a BNE in which the value of (21) is at least $V_j(\mathbf{x}^*, \mathbf{t}^0)$. By (21) and (22), the auction may be interpreted as follows. An asset of size Z is owned in equal shares by a collective of $2n$ agents. Each agent j is risk-neutral and the value he assigns to the asset (or any fraction of it) is $1-\theta_j$, where θ_j is independently drawn from a *cdf* F. In order to decide which of the agent gets full ownership rights of the asset, the agents play a first-price auction. The revenues from the auction are distributed equally among all $2n$ agents. By Proposition 6 in CGK, this auction has an BNE satisfying the participation constraint. ■