A Mechanism-Design Approach to Speculative Trade*

Kfir Eliaz† and Ran Spiegler‡

First version: November 2005
This vesion: July 2006

Abstract

When two agents hold different priors over an unverifiable state of nature, which affects the outcome of a game they are about to play, they have an incentive to bet on the game's outcome. We pose the following question: what are the limits to the agents' ability to realize gains from such speculative bets when their priors are private information? We apply a "mechanism design" approach to this question. We characterize interim-efficient bets and discuss their implementability in terms of the underlying game's payoff structure. In particular, we show that as the costs of unilaterally manipulating the bet's outcome become more symmetric across states and agents, implementation becomes easier.

1 Introduction

A primary task of the mechanism-design literature has been to draw the barriers to trade due to asymmetric information. 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

*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

differences in *beliefs*. What are the limits to 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 neglected these questions, probably because of 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. However, we assume that the agents' priors are private information. We apply a *"mechanism-design approach"* in order to examine the extent to which this form of asymmetric information creates a barrier to speculative bets.

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) provides 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 heterogeneous priors, and show that incentive-compatibility constraints exclude bets as part of the revenue-maximizing mechanism.

We demonstrate our approach with a simple two-period model. In period 2, a pair of agents plays a game whose payoffs depend on an unverifiable state of Nature. The state is commonly known in period 2. However, in period 1 it is unknown to the agents, who hold different prior beliefs over the state and therefore might benefit from betting on it. Since the state is unverifiable, the agents cannot bet on its realization. The set of verifiable contingencies is captured by a partition over the set of action profiles in the game. A bet signed in period 1 is a function that assigns a budget-balanced transfer to each cell in the partition. The agents' priors are private information, but it is common knowledge that they are independently drawn from some distribution F. We define a notion of a "constrained interim-efficient" bet and ask whether it can be implemented in Bayesian equilibrium by some mechanism.

An important feature of this model is that the outcome of a bet can be manipulated by the agents, through their choice of action in period 2. Thus, 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. A constrained interim-efficient bet in such an environment maximizes these gains (formally, the sum of the agents' interim expected utilities, calculated according to their own priors), subject to the constraint that neither agent wishes 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, from a methodological point of view, quasi-linear utility and unbounded transfers are standard assumptions in the mechanism design literature. Second, there are many real-life situations in which agents with heterogenous beliefs bet on outcomes they can manipulate. For instance, a contract signed between an investor and an entrepreneur may reflect their different degrees of optimism regarding the future success of their business venture. At the same time, the contract affects the entrepreneur's incentives to undertake risky projects. Gambling over the outcome of sporting contests provides another example. Different prior beliefs regarding the contest's outcome provide a motive to bet. However, when the stakes are high, a contestant may deliberately tilt the score in order to win a side bet. Finally, when market speculators trade in financial derivatives, they may be driven by different prior beliefs over market fundamentals that affect future stock prices. When the stock market is imperfectly competitive, traders are able to manipulate stock prices, and this affects the positions that they take in the derivative market.

Third, the bounds on the stakes of bets in our model are *endogenous*. This allows us to establish a link between the implementability of constrained interim-efficient bets and the payoff structure of the underlying game. The main result in the paper is that when a constrained interim-efficient bet is "purely speculative" (in the sense that it does not affect the game's outcome), it can be implemented for a larger set of distributions F when the costs of unilateral manipulation of the bet's outcome become more symmetric across states and agents.

The technical basis for this result is a formal analogy to a more conventional mechanism-design model due to Cramton, Gibbons and Klemperer (1987) - CGK henceforth - which extends 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 the usefulness of this analogy with a pair of applications, in which agents are able to bet on the market price that results from some market interaction in which they take part.

2 An example: betting on an agent's future action

Before presenting our model in generality, we wish to convey some of its main ideas through a simple special case. Consider an agent who faces a choice between two actions: a or b. His payoff from each action 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. The payoffs are given by the following table:

	a	b
u	A	C
v	D	B

where $A \geq C$ and $B \geq D$, 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 party, 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 two parties have different priors, they 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 parties can bet on the agent's action, which is 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* t is a function that assigns a pair of monetary transfers, $t_1(x)$ and $t_2(x)$, to every $x \in \{a, b\}$, where $t_i(x)$ is the amount that party i receives if the agent chooses x in period 2. The transfers are budget-balanced - i.e., $t_1(x) = -t_2(x)$. If the parties 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 $u(x) + t_2(x)$ in state u and $v(x) + t_2(x)$ in state v, and the speculator's utility from the agent's action x is $t_1(x)$, regardless of the state. If no bet is signed, the agent faces the "bare" decision problem and the speculator receives nothing.

This example fits a number of real-life situations. The parties can be interpreted as a buyer and a seller. In period 1, the buyer does not know which of two varieties of a product will fit his needs in period 2. A bet is essentially an advance contract which specifies a price for each variety (if no deal is signed in period 1, the buyer purchases the product from an alternative supplier). 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 state of the economy. 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 state of the economy.

Consider a bet t, and suppose that both parties 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 $\theta_1 \cdot t_1(x^u) + (1 - \theta_1) \cdot t_1(x^v)$, while the agent's is $\theta_2 \cdot [u(x^u) - t_1(x^u)] + (1 - \theta_2) \cdot [v(x^v) - t_1(x^v)]$. The term "*interim*" is fitting because it refers to the parties' expected payoffs upon learning their prior. The sum of the parties' interim expected payoffs can be conveniently written as

$$\theta_2 \cdot u(x^u) + (1 - \theta_2) \cdot v(x^v) + (\theta_1 - \theta_2) \cdot [t_1(x^u) - t_1(x^v)] \tag{1}$$

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 parties would want to sign: if $\theta_1 > \theta_2$, they would set $t_1(x^u) \gg t_1(x^v)$, and if $\theta_1 < \theta_2$, they would set $t_1(x^v) \gg t_1(x^u)$. However, because the agent cannot commit to his second-period action, the parties must take into account his ability to manipulate the bet's outcome. For instance, suppose that the parties agree on a bet that satisfies $t_1(b) - t_1(a) > B - D$. 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 parties cannot benefit from betting on the agent's action. *Thus, in order to be sustainable, a non-trivial bet must provide the agent with incentives to take different actions in different states.*[1]

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

$$\begin{aligned} u(x^u) - t_1(x^u) &\geq u(x^v) - t_1(x^v) \\ v(x^v) - t_1(x^v) &\geq v(x^u) - t_1(x^u) \end{aligned}$$

which we call "second-period incentive compatibility" (**SPIC**) constraints. The efficiency criterion employed here is standard Pareto efficiency: for any non-CIE (x, t),

[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.

there is another pair (x', t') which satisfies the SPIC constraints and yields higher interim expected utility for both parties. If (x, t) is a solution to the constrained optimization problem, we refer to t as a *CIE bet*. We refer to the expression (1), evaluated at a CIE pair (x, t), as the *CIE surplus*.

It follows from (1) that if $\theta_1 > \theta_2$, the parties would want to set $t_1(x^u) - t_1(x^v)$ 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 $t_1(x^u) - t_1(x^v)$ to be equal to the *lower* bound implied by the SPIC constraints, $v(x^u) - v(x^v)$. Both bounds on $t_1(x^u) - t_1(x^v)$ are relaxed to their utmost when $x^u = a$ and $x^v = b$. Thus, we have the following characterization.

Remark 1 *A pair (x, t) is CIE 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:

$$t_1(a) - t_1(b) = \begin{cases} A - C & if \quad \theta_1 > \theta_2 \\ D - B & if \quad \theta_1 < \theta_2 \end{cases}$$

We now turn to the question of whether the CIE surplus can be implemented when the parties' priors are not common knowledge. We assume that each party privately and independently draws his prior on u 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, let $A - C = B - D$, and suppose that in period 1, the parties play the following naïve mechanism: each party guesses the agent's second-period action; when *exactly* one party guesses correctly, he receives $(A - C)/2$ from the other party; otherwise, no payments are made in period 2. Since $\theta_1 \neq \theta_2$ with probability one, the two parties can always earn speculative gains, if the party with the higher prior on u guesses a while the other guesses b. However, note that when $\theta_1, \theta_2 > \frac{1}{2}$, both parties would want to guess a. Similarly, when $\theta_1, \theta_2 < \frac{1}{2}$, both parties would want to guess b. Consequently, for this range of (θ_1, θ_2), the fact that the parties' priors are private information implies that they forgo potential speculative gains.

We consider the problem of implementing the CIE surplus via a *direct* mechanism. This means that the parties play a two-period game, denoted Γ. In the first period, each party i submits a report $\hat{\theta}_i \in [0, 1]$ (interpreted as his stated prior on u), or chooses not to participate. If at least one party chooses the latter, the agent faces the "bare" decision problem. If both parties choose to participate, every pair of reports $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \hat{\theta}_2)$ is assigned a bet $t(\hat{\boldsymbol{\theta}})$, which is disclosed to the agent. In period 2, after the state of Nature is realized, the agent chooses an action x and pays $t_1(x \mid \hat{\boldsymbol{\theta}})$. In state u, he chooses x to maximize $u(x) - t_1(x \mid \hat{\boldsymbol{\theta}})$, whereas in state v, he chooses x to maximize $v(x) - t_1(x \mid \hat{\boldsymbol{\theta}})$.

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

Proposition 1 *There exists a distribution F for which the CIE surplus 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 CIE surplus is implementable expands. When $A - C = B - D$, the CIE surplus is implementable for every distribution F.*[2]

The main lesson from this result is that implementability of the CIE surplus diminishes as the agent's incentives to manipulate the outcome of the CIE bet become more uneven across states. To develop an intuition for this result, consider a mechanism that satisfies $t_1(a \mid \hat{\boldsymbol{\theta}}) - t_1(b \mid \hat{\boldsymbol{\theta}}) = A - C$ if $\hat{\theta}_1 \geq \hat{\theta}_2$ and $t_1(a \mid \hat{\boldsymbol{\theta}}) - t_1(b \mid \hat{\boldsymbol{\theta}}) = D - B$ if $\hat{\theta}_1 < \hat{\theta}_2$. Then, regardless of the first-period outcome, the agent takes the ex-post efficient action in each state. Moreover, if $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$, the bet t is CIE. The problem is to design such a mechanism $t(\hat{\boldsymbol{\theta}})$, which also ensures that the parties participate and report their true priors.

Our approach to analyzing this problem involves reinterpreting it as a problem of allocating an asset to the person who values it the most. Suppose that both parties report their true priors in period 1 and consider the agent's decision problem in period 2. What is his gain from choosing the efficient action relative to choosing b? By definition, the gain is zero in state v, regardless of whether θ_1 is higher or lower than θ_2. However, in state u the gain is $(A - C) - [t_1(a \mid \boldsymbol{\theta}) - t_1(b \mid \boldsymbol{\theta})]$. By our construction of $t(\hat{\boldsymbol{\theta}})$ and the assumption that $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$, this difference is equal to zero when $\theta_1 \geq \theta_2$ and equal to $(A - C) + (B - D)$ when $\theta_1 < \theta_2$.

Thus, the agent's gain may be interpreted as a right to receive a prize of $(A - C) + (B - D)$ conditional on choosing a in period 2. Put differently, the right is an asset of size $(A - C) + (B - D)$, whose first-period valuation by each party i is $\theta_i \cdot [(A - C) + (B - D)]$. Note that the agent receives this asset if and only if $\theta_1 < \theta_2$.

[2] This result is a special case of Corollary 1 (see Section 4). Therefore, we do not provide a specific proof for it. All proofs appear in the Appendix.

This is analogous to allocating the asset to the party who values it the most. What happens when no bet is signed in period 1? The agent's gain from choosing the efficient action relative to choosing b is zero in state v and $A - C$ in state u. Thus, it is as if the agent initially holds a share of $A - C$ in the asset described above. His first-period valuation of this asset is $\theta_2(A - C)$. By signing the bet with the speculator, the agent increases his share by $B - D$, as long as $\theta_2 > \theta_1$.

These observations suggest that the problem of implementing the CIE surplus is analogous to the problem of dissolving a partnership efficiently. In this problem, two parties jointly hold an asset of size $(A - C) + (B - D)$. The parties' shares in the asset are $A - C$ and $B - D$. Each party privately and independently draws a valuation of the asset. The problem is to design a mechanism that allocates the entire asset to the party with the highest valuation, subject to the constraint that both parties agree to participate in this mechanism.

CGK showed that implementing this objective depends on the initial ownership structure. When $A - C \gg B - D$ - that is, if the agent enters the negotiation mostly a "seller" of the asset - the same forces that underlie the Myerson-Satterthwaite theorem make it hard to allocate the asset efficiently. As the gap between $A - C$ and $B - D$ shrinks, each party enters the negotiation both as a seller and a buyer, and thus he has "countervailing incentives" when reporting his valuation. Translated into the language of our model, this result means that implementing the CIE bet becomes easier when the agent's costs of unilaterally manipulating the bet become more equal across states.

3 The model

A *bilateral speculation problem* has the following components. There are two periods. In period 2 a pair of agents, $i = 1, 2$, play a normal form game with complete information denoted by G. The set of actions available to agent i is denoted A_i. A partition X is defined on the set of action profiles $A_1 \times A_2$, such that $x(a_1, a_2)$ denotes the cell in the partition that contains the action profile (a_1, a_2). We interpret X as the set of "verifiable outcomes". For example, when G represents a sports competition, a cell in X may consist of all action profiles which induce a particular final score. When G is a market game, a cell in X may consist of all action profiles which induce a particular trading price.

The payoffs in G depend on the state of Nature, which is common knowledge in period 2. There are two possible states, u and v. In one state, player i's utility function from each action pair is $u_i : A_1 \times A_2 \to \mathbb{R}$, while in the second state this function is

$v_i : A_1 \times A_2 \to \mathbb{R}$. Let $G(\omega)$ denote the second-period game played in state $\omega \in \{u, v\}$. We assume that $G(\omega)$ has a pure-strategy NE for every state ω. Let a^ω denote the action profile that is played in state ω. Denote $x^\omega = x(a^\omega)$.

In period 1, before the state is realized, the two agents hold different prior beliefs over the states of Nature: agent i assigns probability θ_i to state u. These are purely differences in *prior* opinions. This means that if agent i knew θ_j, this would not cause him to update his belief regarding the state of Nature. Each agent independently and privately draws his prior from the same, commonly known continuous *cdf* on $[0, 1]$, denoted F. We represent a bilateral speculation problem by the tuple $\langle (u, v), G, X, F \rangle$.

A bet t is a function that assigns a pair of budget-balanced transfers (t_1, t_2) to every cell in X. Let $t_i[x(a_1, a_2)]$ denote the transfer that agent i receives from agent j, when the action pair (a_1, a_2) is played. Budget-balancedness means that $t_1[x(a_1, a_2)] = -t_2[x(a_1, a_2)]$. Given a bet t, the payoff of agent i from the action profile $a = (a_1, a_2)$ is $u_i(a) + t_i[x(a)]$ in state u and $v_i(a) + t_i[x(a)]$ in state v. We refer to the second-period game induced by a bet t as the "modified game", and denote it by $G(\omega, t)$.

Definition 1 *A triple (a^u, a^v, t) is constrained interim-efficient (CIE) for a given pair of priors (θ_1, θ_2), if it solves the optimization problem*

$$\max_{(a^u, a^v, t)} \sum_{i=1}^{2} \{\theta_i[u_i(a^u) + t_i(x(a^u))] + (1 - \theta_i)[v_i(a^v) + t_i(x(a^v))]\} \tag{2}$$

subject to the constraints

$$\begin{aligned} u_i(a_i^u, a_j^u) + t_i[x(a_i^u, a_j^u)] &\geq u_i(a_i', a_j^u) + t_i[x(a_i', a_j^u)] \\ v_i(a_i^v, a_j^v) + t_i[x(a_i^v, a_j^v)] &\geq v_i(a_i', a_j^v) + t_i[x(a_i', a_j^v)] \end{aligned} \tag{SPIC}$$

for $i = 1, 2$ and for all $a_i' \in A_i$.

As in the example of Section 2, we refer to the value of the objective function (2), evaluated at a CIE tuple (a^u, a^v, t), as the CIE surplus. We refer to a bet t as CIE if there exist action profiles a^u and a^v such that (a^u, a^v, t) is CIE.[3]

[3]Note that if the modified game $G(\omega, t)$ does not have a pure-strategy NE, then t is ruled out as far as constrained interim-efficiency is concerned. Since the bare game is assumed to have a pure-strategy NE in each state, the constrained optimization problem has a solution.

Proposition 2 *A bilateral speculation problem $\langle (u,v), G, X, F\rangle$ with bounded u and v has a finite CIE surplus.*

Thus, as long as u and v are bounded, the CIE surplus is well-defined - that is, the infinite-bets problem does not exist in our model.

Discussion of our epistemic assumptions

A key ingredient in our model is the assumption that the agents' conflicting beliefs are due to heterogeneous prior opinions. In particular, their beliefs cannot be derived from a common prior via Bayes' rule. To see why, assume that the agents shared a common prior belief, where $p(\omega,\theta_1,\theta_2)$ denotes the prior probability that the state of Nature is ω, agent 1's type is θ_1 and agent 2's type is θ_2. The posterior probability that type θ_i of agent i assigns to u is θ_i. Our assumption that knowing the opponent's type does not cause an agent to update his beliefs regarding the state of Nature implies that for every $\theta_i,\theta_j,\theta'_j$, $p_i(u \mid \theta_i,\theta_j) = p_i(u \mid \theta_i,\theta'_j)$. That is:

$$\frac{p(u,\theta_i,\theta_j)}{p(u,\theta_i,\theta_j)+p(v,\theta_i,\theta_j)} = \frac{p(u,\theta_i,\theta'_j)}{p(u,\theta_i,\theta'_j)+p(v,\theta_i,\theta'_j)}$$

But since agent j's belief regarding the state of Nature is unaffected by knowledge of θ_i, the L.H.S and R.H.S of this equation are the posterior probabilities that types θ_j and θ'_j assign to u. Thus, $\theta_j = \theta'_j$, a contradiction.

The assumption that F is common knowledge is made mainly for methodological reasons, since we wish to parallel the simplest textbook mechanism-design models. One interpretation of this assumption is that in many instances, θ_i is best viewed as agent i's *degree of optimism*. For instance, when G is a price-competition game, u may be characterized by a lower cost of production than v. Alternatively, when G is a bilateral-trade game, there may be larger gains from trade in u than in v. Optimism is a personal trait which is as characteristic of an individual as his valuation of a tradable object in a standard model. Thus, the question of whether F is common knowledge is as pertinent to our model as it is to standard models of trade based on differences in tastes.[4]

An alternative interpretation is that there is a distribution of prior opinions in the general population. Agents become familiar with this distribution by observing a

[4]However, Yildiz (2004) argues that the there is a tension between equilibrium analysis and the interpretation of an agent's prior over states of nature as reflecting his degree of optimism. The reason is that it is unclear why the agent's optimism does not extend to the formation of beliefs regarding the opponent's strategy.

public poll. The common-knowledge assumption means that all agents share the same beliefs regarding the poll's accuracy.

3.1 Purely speculative CIE bets

Since bets are essentially side transfers that modify the payoffs of the second-period game, they can be used not only for speculation, but also as means for sustaining collusion. The speculative role of bets can best be isolated when the agents attain the CIE surplus with a bet that does not affect their second-period behavior, in the sense that their choice of actions is the same as in the absence of bets. Such a CIE bet may be viewed as "purely speculative", since it serves purely as a means for realizing speculative gains.

Definition 2 *We say that the CIE surplus is attained by pure speculation if there exists a pair of action profiles, (a^u, a^v), with the following properties: (i) a^u and a^v are Nash equilibria in $G(u)$ and $G(v)$ respectively, and (ii) for every pair of priors $\boldsymbol{\theta}$, there exists a bet $t(\boldsymbol{\theta})$ such that $[a^u, a^v, t(\boldsymbol{\theta})]$ is CIE for $\boldsymbol{\theta}$. In this case, we say that $[a^u, a^v, t(\boldsymbol{\theta})]$ is a purely speculative CIE tuple and that $t(\boldsymbol{\theta})$ is a purely speculative CIE bet.*

Note that in general, attaining the CIE surplus may require a^u and a^v to vary with $\boldsymbol{\theta}$. However, when the CIE surplus is attained with a purely speculative bet, a^u and a^v are not only independent of the priors, but also a NE of the bare game.

In order to characterize purely speculative CIE bets, we shall need the following notation. For each verifiable outcome x and for each action $a_j \in A_j$, define $A_i(x, a_j) \equiv \{a_i \in A_i : x(a_i, a_j) = x\}$. That is, $A_i(x, a_j)$ is the (possibly empty) set of actions for agent i that induce the verifiable outcome x whenever agent j plays a_j. Let $d_i(a^\omega \to x)$ be the minimal cost that agent i incurs (in terms of his bare-game payoff) when he unilaterally changes the outcome of $G(\omega)$, from $a^\omega = (a_i^\omega, a_j^\omega)$ to an action profile that belongs to the verifiable outcome $x \in X$. Formally:

$$d_i(a^u \to x) \equiv \begin{cases} \min_{a_i' \in A_i(x, a_j^u)} [u_i(a_i^u, a_j^u) - u_i(a_i', a_j^u)] & if \quad A_i(x, a_j^u) \neq \varnothing \\ \infty & if \quad A_i(x, a_j^u) = \varnothing \end{cases}$$

Define $d_i(a^v \to x)$ in a similar manner. Note first that if (a^u, a^v, t) is a CIE tuple, then $d_i(a^\omega \to x(a^\omega)) = 0$, because t is constant over all action profiles in x. If $d_i(a^\omega \to x(a^\omega)) < 0$, then agent i would have a profitable deviation from a_i^ω, in contradiction to a^ω being a NE of $G(\omega, t)$. Also note that if $[a^u, a^v, t(\boldsymbol{\theta})]$ is a *purely speculative* CIE tuple, then $d_i(a^\omega \to x) \geq 0$, because a^ω is a NE of $G(\omega)$.

For the final piece of notation, let

$$\begin{aligned} D_1(a^u, a^v) &\equiv \min_{y \in X} \quad d_1(a^u \to y) + d_2(a^v \to y) \\ D_2(a^u, a^v) &\equiv \min_{y \in X} \quad d_2(a^u \to y) + d_1(a^v \to y) \end{aligned}$$

$$D^*(a^u, a^v \quad | \quad \boldsymbol{\theta}) = \begin{cases} D_2(a^u, a^v) & if \quad \theta_1 \geq \theta_2 \\ -D_1(a^u, a^v) & if \quad \theta_1 < \theta_2 \end{cases}$$

When G is not finite, we assume that u and v are such that $D_1(a^u, a^v)$ and $D_2(a^u, a^v)$ are well-defined.

Proposition 3 *A purely speculative CIE bet t satisfies the following property for every pair of priors $\boldsymbol{\theta}$:*

$$t_1(x^u) - t_1(x^v) = D^*(a^u, a^v \mid \boldsymbol{\theta}) \tag{3}$$

Thus, the stakes of purely speculative CIE bets are determined by how costly it is for agents (in terms of bare-game payoffs) to manipulate the bet's outcome unilaterally. To understand the meaning of $D^*(a^u, a^v \mid \boldsymbol{\theta})$, consider two special cases. First, recall the example of Section 2, in which one of the agents, agent 1, has a degenerate action set, while the opponent has a pair of available actions. In this case, $D_1(a^u, a^v) = v_2(a^v) - v_2(a^u)$ and $D_2(a^u, a^v) = u_2(a^u) - u_2(a^v)$. The stakes of the purely speculative CIE bet are thus determined by agent 2's costs of unilaterally manipulating the bet's outcome from a^u into a^v in state u, and from a^v into a^u in state v.

Second, consider a symmetric Bertrand model, in which the firms' marginal cost in state ω is c^ω, $\omega \in \{L, H\}$, $c^L < c^H$. Assume that a verifiable outcome is the market price induced by the firms' bids. We analyze this example in detail in Section 4. In particular, we show that the CIE surplus can be sustained if firms play the bare-game NE in each state. While neither firm can manipulate the market price in state L upward, each firm can manipulate the market price in state H, from c^H downwards to c^L. Other market prices turn out not to matter. Thus, the stakes of the purely speculative CIE bet are determined by the two firms' cost of unilaterally manipulating the bet's outcome from x^H into x^L. Specifically, $D_1(a^u, a^v) = D_2(a^u, a^v) = c^H - c^L$.

In more complicated situations, we also need to take into account manipulation of the bet's outcome from a^ω into an outcome y which never occurs in any state in equilibrium. To see the origin of the expression for $D_1(a^u, a^v)$ in this more general case, suppose that agent 1 has bet against x^u, presumably because he thought that u was unlikely. Now, when the state u occurs and the outcome x^u is expected to be realized, agent 1 may wish to manipulate the bet's outcome. One possibility is to impose an outcome in x^v, in which case agent 1 suffers a bare-game loss of $d_1(a^u \to x^v)$. Clearly, the side-bet difference $t_1(x^v) - t_1(x^u)$ cannot exceed this amount. But another way is to impose an outcome $y \neq x^v$, in which case agent 1 suffers a bare-game loss of $d_1(x^u \to y)$. By budget-balancedness, this affects the bounds on $t_1(x^v) - t_1(x^u)$, through the possibility that agent 2 will manipulate the bet's outcome from a^v to y.

3.2 Implementation of purely speculative CIE bets

As in Section 2, we study implementation by a direct mechanism. In period 1, each agent i submits a report $\hat{\theta}_i \in [0,1]$ or chooses not to participate. If at least one agent chooses the latter, the agents play $G(\omega)$ in state ω. Otherwise, every profile of reports $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \hat{\theta}_2)$ is assigned a bet $t(\hat{\boldsymbol{\theta}})$, and the agents play $G(\omega, t(\hat{\boldsymbol{\theta}}))$ in period 2. Thus, a direct mechanism $t(\hat{\boldsymbol{\theta}})$ induces a two-stage game with incomplete information, denoted $\Gamma(t)$.

Define $T_i^u(\theta_i') \equiv E_{\theta_j} t_i(x^u \mid \theta_i', \theta_j)$ and $T_i^v(\theta_i') \equiv E_{\theta_j} t_i(x^v \mid \theta_i', \theta_j)$. That is, if agent i reports a prior θ_i', while agent j is truthful, then $T_i^\omega(\theta_i')$ is agent i's expected transfer in state ω under the mechanism $t(\boldsymbol{\theta})$.

Definition 3 *Suppose that the CIE surplus is attained by pure speculation. A direct mechanism $t(\hat{\boldsymbol{\theta}})$ implements the CIE surplus for a given distribution F if:*

(EFF) $t(\hat{\boldsymbol{\theta}})$ satisfies (3),

and there exists a PBNE in $\Gamma(t)$ satisfying:

(PS-SPIC) The second-period action profile in state ω is a^ω after every history, where a^ω is a pure-strategy NE in $G(\omega)$.

(IC) Each agent reports his true prior in period 1, conditional on participating. That is, for every $i = 1,2$ and every θ_i, θ_i':

$$\theta_i[T_i^u(\theta_i) - T_i^v(\theta_i)] + T_i^v(\theta_i) \geq \theta_i[T_i^u(\theta_i') - T_i^v(\theta_i')] + T_i^v(\theta_i')$$

(IR) Each agent chooses to participate in period 1. That is, for every $i = 1,2$ and every θ_i:

$$\theta_i[T_i^u(\theta_i) - T_i^v(\theta_i)] + T_i^v(\theta_i) \geq 0$$

The EFF condition means that if the agents report truthfully, then $t(\hat{\boldsymbol{\theta}})$ is a CIE bet. Condition PS-SPIC means that in the second stage of $\Gamma(t)$, the agents play a NE of the bare game, independently of the first-stage outcome. This means that we are forcing the mechanism to be purely speculative. The IC and IR constraints refer to the agents' first-period decisions. Note that because of the pure speculation assumption, these constraints suppress any reference to the bare-game payoffs, .

Our goal is to establish a relation between implementation of the pure speculation CIE surplus in a bilateral speculation problem and implementation of efficient dissolution of a partnership. This latter problem is defined as follows. A two-member partnership is a triple $\langle r_1, r_2, F\rangle$, where r_i is partner i's initial share in the jointly owned asset and F is the continuous distribution on $[0,1]$ from which both partners independently (but privately) draw their valuations of the asset. The partners are assumed to be risk neutral with quasi-linear preferences, where θ_i denotes partner i's value for a unit of the asset. A partnership is dissolved efficiently if the entire asset $r_1 + r_2$ is allocated to the partner with the highest valuation.

A direct mechanism for dissolving a partnership is a pair of functions $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ that assign, for each pair of reported values $\hat{\boldsymbol{\theta}}$, an allocation of shares, $q_1(\hat{\boldsymbol{\theta}})$ and $q_2(\hat{\boldsymbol{\theta}})$, and a pair of monetary transfers, $m_1(\hat{\boldsymbol{\theta}})$ and $m_2(\hat{\boldsymbol{\theta}})$, such that for all $\hat{\boldsymbol{\theta}}$, $q_i(\hat{\boldsymbol{\theta}}) \geq 0$, $q_1(\hat{\boldsymbol{\theta}}) + q_2(\hat{\boldsymbol{\theta}}) = r_1 + r_2$ and $m_1(\hat{\boldsymbol{\theta}}) + m_2(\hat{\boldsymbol{\theta}}) = 0$.

Definition 4 *A mechanism $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ efficiently dissolves a partnership $\langle r_1, r_2, F\rangle$ if it satisfies the following properties for $i = 1, 2$:*

(EFF) Whenever $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$,*

$$q_i(\boldsymbol{\theta}) = \begin{cases} r_1 + r_2 & if \quad \theta_1 \geq \theta_2 \\ 0 & if \quad \theta_1 < \theta_2 \end{cases}$$

(IC) There is a Bayesian NE in which every partner reports his true value. That is, for every $i = 1, 2$ and every θ_i, θ_i':*

$$\theta_i Q_i(\theta_i) + M_i(\theta_i) \geq \theta_i Q_i(\theta_i') + M_i(\theta_i')$$

where $Q_i(\hat{\theta}_i) \equiv E_{\theta_j} q_i(\hat{\theta}_i, \theta_j)$ and $M_i(\hat{\theta}_i) \equiv E_{\theta_j} m_i(\hat{\theta}_i, \theta_j)$.

(IR) Each partner's interim-expected payoff in the truth-telling Bayesian NE is at least as high as the value he assigns to his initial share. That is, for every $i = 1, 2$ and every*

θ_i :

$$\theta_i Q_i(\theta_i) + M_i(\theta_i) \geq \theta_i r_i$$

We say that a partnership can be dissolved efficiently if there exists a direct mechanism that implements its efficient dissolution. We are now ready for the main result of this paper.

Proposition 4 *Let $\langle (u,v), G, X, F \rangle$ be a bilateral speculation problem with a CIE surplus that is attained by pure speculation and sustains a^{ω} in state ω. The CIE surplus is implementable for F if and only if the partnership $\langle D_1(a^u, a^v), D_2(a^u, a^v), F \rangle$ can be efficiently dissolved.*

Thus, implementing a pure speculation CIE surplus is equivalent to implementing efficient dissolution of a partnership, where the size of the jointly owned asset is $D_1(a^u, a^v) + D_2(a^u, a^v)$, and the partners' shares are $D_1(a^u, a^v)$ and $D_2(a^u, a^v)$. We can therefore utilize Propositions 1-3 in CGK, and obtain the following corollary. Let

$$\rho = \frac{D_1(a^u, a^v)}{D_1(a^u, a^v) + D_2(a^u, a^v)}$$

Corollary 1 *Suppose that the bilateral speculation problem $\langle (u,v), G, X, F \rangle$ has a CIE surplus that is attained by pure speculation. Then, there exists a distribution F for which the CIE surplus is implementable, if and only if $\rho \in (0,1)$. Moreover, as ρ becomes closer to $\frac{1}{2}$, the set of such distributions F expands. When $\rho = \frac{1}{2}$, the CIE surplus is implementable for every F.*

To see the meaning of this result, suppose that we can ignore the possibility that agents manipulate the bet's outcome into some $y \neq x^u, x^v$. In this case:

$$\begin{aligned} D_1(a^u, a^v) &\equiv \min \quad \{d_1(a^u \to x^v), d_2(a^v \to x^u)\} \\ D_2(a^u, a^v) &\equiv \min \quad \{d_2(a^u \to x^v), d_1(a^v \to x^u)\} \end{aligned}$$

This means that implementability of the CIE surplus depends on either: *(i)* the extent to which the costs of manipulating the bet in one of the states are asymmetric across agents; or *(ii)* the extent to which the costs of manipulating the bet for one of the agents are asymmetric across states. As these asymmetries vanish, the set of distributions F for which the CIE surplus is implementable expands.[5]

When $G(u)$ and $G(v)$ are symmetric games, and a^u and a^v are symmetric NE in $G(u)$ and $G(v)$, we have $D_1(a^u, a^v) = D_2(a^u, a^v)$. In this case, our implementation problem is equivalent to the equal-share partnership dissolution problem, which CGK show to be implementable for any F. Thus, symmetric speculation problems occupy a special place in our model.

4 Applications

In this section we apply the main result to environments in which agents play a market game in period 2, and bet on its outcome in period 1.

4.1 Bertrand competition

In this sub-section, the second-period bare game G is a standard Bertrand competition, where each seller $i \in \{1, 2\}$ chooses a price $a_i \in \mathbb{R}$ (we allow for negative prices). The market price induced by a^ω is $p^\omega = \min(a_1^\omega, a_2^\omega)$. The sellers have identical marginal costs, which are fixed and may be $c^H > 0$ in state H or $c^L \in [0, c^H)$ in state L. Let θ_i denote seller i's prior on L. In period 1, the sellers can sign a bet that is contingent only on the second-period market price. Thus, $x(a_1, a_2) = \min\{a_1, a_2\}$.

Proposition 5 *In the above bilateral speculation problem, the CIE surplus is attained by pure speculation. Moreover:*

(i) The CIE surplus is sustained by a triple (a^L, a^H, t) such that:

$$\begin{aligned} a^\omega &= (c^\omega, c^\omega) \text{ for every } \omega = L, H \\ t_i(p) &= t_i(c^H) \text{ for all } p > c^L \\ t_i(p) &= t_i(c^H) + c^H - c^L \text{ for all } p \leq c^L \end{aligned}$$

where $i = \arg\max(\theta_1, \theta_2)$.

(ii) The CIE surplus is $|\theta_1 - \theta_2| \cdot (c^H - c^L)$.

[5] When deviations to outcomes $y \neq x^u, x^v$ cannot be ignored, the exact form of asymmetry which is relevant for the corollary is somewhat harder to interpret.

Under the purely speculative CIE bet, both sellers play $a^\omega = c^\omega$ in each state ω. Therefore, their bare-game payoff is zero, and their interim surplus is derived from the side bets only. The stakes of their bet are determined by the cost of unilaterally lowering the price in state H, from c^H to c^L. In contrast, no seller can unilaterally manipulate the market price in state L upward.

The proof of this result is not trivial. Pure speculation implies $p^H = c^H$ and $p^L = c^L$. One could imagine that if we extended the gap between p^H and p^L, we might be able to relax the SPIC constraints and thereby increase the stakes of CIE bets. However, we show that in order for this to be sustainable, there must be a state ω for which $p^\omega < c^\omega$. The challenging part in the proof is to show that the SPIC constraints that are required in order to sustain a price below the marginal cost are too stringent.

Corollary 2 *In the above bilateral speculation problem, the CIE surplus is implementable for every F.*

The reason for this result is that $D_1(a^H, a^L) = D_2(a^H, a^L) = c^H - c^L$, and by Corollary 1, our implementation problem is equivalent to an equal-share partnership dissolution problem.

4.2 Bilateral trade

In this sub-section, the second-period bare game involves bilateral trade. A seller, denoted s, owns one unit of an indivisible good. The value of the good to the seller is c. A potential buyer, denoted b, evaluates the good at l or h, where $h > c > l$. In period 2, when the buyer's valuation becomes common knowledge, the two agents play a double auction: they simultaneously submit ask and bid prices, p_s and p_b; if $p_b \geq p_s$, trade takes place at a price $\frac{1}{2}p_b + \frac{1}{2}p_s$; and if $p_b < p_s$, there is no trade. Thus, if there is trade at a price p when the buyer's valuation is $\omega \in \{l, h\}$, then the buyer's payoff is $\omega - p$ and the seller's payoff is $p - c$. If there is no trade, both agents earn a payoff of zero. We allow bid and ask prices to be arbitrarily positive or arbitrarily negative.

We assume that the agents can only bet on whether trade takes place, and at what price. Thus, if (p_b, p_s) and (p'_b, p'_s) induce the same market price, or if both result in no trade, then $x(p_b, p_s) = x(p'_b, p'_s)$. We use the following abbreviated notation. If (p_b, p_s) induces trade at a price p, we write $x = p$. If (p_b, p_s) induces no trade, we write $x = NT$. Let θ_b and θ_s denote the prior probabilities that the buyer and seller assign to h.

Proposition 6 *In the above bilateral speculation problem, the CIE surplus is attained by pure speculation. Moreover:*

(i) The value of the CIE surplus is

$$\max(\theta_s, \theta_b) \cdot (h - c)$$

(ii) The CIE surplus is sustained by any (a^l, a^h, t) *for which:*

$$\begin{aligned}
p_s^l &\gg 0,\ p_b^l \ll 0 \ \textit{(hence } x^l = NT\textit{)} \\
p_s^h &= p_b^h = \frac{h+c}{2} \\
t_s(x) &= \begin{cases} t_s(NT) + \frac{h-c}{2} & if \quad \theta_s > \theta_b \\ t_s(NT) - \frac{h-c}{2} & if \quad \theta_s < \theta_b \end{cases} \ \textit{for any } x \neq NT
\end{aligned}$$

Observe that the CIE bet conditions only on whether trade takes place, and does not distinguish between different trading prices.

Corollary 3 *The CIE surplus is implementable for every* F*.*

To see why this corollary holds, note that the action profiles (p_s^l, p_b^l) and (p_s^h, p_b^l) are bare-game NE in states l and h, respectively. Also note that part *(ii)* in Proposition 6 implies $D_1(a^l, a^h, t) = D_2(a^l, a^h, t) = \frac{h-c}{2}$. Therefore, by Corollary 1, the CIE surplus is implementable for any F.

This result relies on a suitable selection of the equilibrium market price in state h. The bare game $G(h)$ has a continuum of NE. It can be shown that for each of these equilibria a^h, there exists a purely speculative CIE bet t. However, these alternative equilibria would imply $D_1(a^l, a^h) \neq D_2(a^l, a^h)$, and therefore we would not be able to claim that implementation is possible for all distributions F. It turns out that there is a unique trading price $p^h = \frac{h+c}{2}$ for which we can construct a tuple (a^h, a^l, t) such that $D_1(a^l, a^h) = D_2(a^l, a^h)$. Thus, the requirement that the CIE surplus be implementable for all F pins down the market price in state h.

5 Discussion

In this section, we discuss extensions and elaborations of our model, as well as related literature.

An indirect mechanism

The purely speculative, CIE tuples (a^u, a^v, t) derived in the applications of Section 4 share two properties. First, $D_1(a^u, a^v) = D_2(a^u, a^v)$. Second, t has only *two* values in its range. In other words, there is a two-cell partition of the set of verifiable outcomes, $\{X^u, X^v\}$, such that $t(x) = t(x^u)$ for every $x \in X^u$, and $t(x) = t(x^v)$ for every $x \in X^v$. In the example of Section 2, the first property holds whenever $A - C = B - D$, while the second property holds automatically because the agent has only two actions.

It can be shown that these properties imply that for any F, the CIE surplus can be implemented by the following *indirect* mechanism. In period 1, the agents play a sealed-bid, first-price auction in which: *(i)* the revenues are equally shared among the bidders; *(ii)* the highest-bidding agent wins the right to receive a transfer of $D_2(a^u, a^v)$ from the other agent if and only if the second-period outcome is in X^u. The proof of this result, which is omitted for the sake of brevity, adapts Propositions 5 and 6 in CGK to the language of our model.

In the Bertrand example of Sub-Section 4.1, this indirect mechanism means that the sellers play a first-price auction for the right to receive a prize of $c^H - c^L$ as long as the market price does not exceed c^L. In the bilateral trade example of Sub-Section 4.2, the two parties play a first-price auction for the right to receive a prize of $\frac{h-c}{2}$ whenever trade occurs. In the example of Section 2, the mechanism means that the parties play a first-price auction in order to determine which of them wins the right to a prize of $A - C$ conditional on the agent choosing a in period 2.

In all three cases, the CIE bet may be interpreted as a future contract (which is essentially a step function of the market price in the Bertrand example, or a function of whether the market clears in the bilateral trade example, or a function of the agent's action in the example of Section 2), competed for in a market which is designed as a first-price auction. Thus, the indirect mechanism may serve as a theoretical benchmark for the design of market institutions for speculative trade in derivatives.

Impurely speculative bets

Our main result concerns the implementability of pure-speculation CIE surplus. We have given a number of examples, in which the CIE bets are indeed purely speculative, and therefore the main result applies. However, in some cases, constrained interim-efficiency is inconsistent with pure speculation: second-period behavior depends on the bet signed in the first period, and therefore on the agents' priors. For instance, modify the bilateral trade example of Sub-Section 4.2 such that $l > c$. The ex-post efficient outcome now involves trade in both states. Using the same methods of derivation as in Sub-Section 4.2, it can be shown that the CIE surplus is

$$\max(\theta_s, \theta_b) \cdot (h - c) + [1 - \min(\theta_s, \theta_b)] \cdot (l - c)$$

and in particular, the market outcome is ex-post efficient in both states.

In order for CIE bets to be purely speculative, the assignment of market prices to states must be independent of the agents' priors. Thus, for every $\omega = l, h$, there must be a trading price p^ω which is independent of (θ_s, θ_b). Denote $p^* = \max(p^l, p^h)$ and $p_* = \min(p^l, p^h)$. If $p^* = p_*$, then total surplus is $\theta_b h + (1 - \theta_b)l - c$, which is below the CIE surplus. Therefore, $p^* > p_*$.

Suppose that $p^* = p^h$ and $p_* = p^l$. The seller can unilaterally lower the price in state h from p^h to p^l. The following SPIC constraint prevents him from doing so:

$$p^h - c + t_s(p^h) \geq p^l - c + t_s(p^l)$$

Therefore, $p^h + t_s(p^h) - p^l - t_s(p^l) \geq 0$. The expression for total surplus is:

$$[\theta_b h + (1 - \theta_b)l - c] + (\theta_s - \theta_b) \cdot [p^h + t_s(p^h) - p^l - t_s(p^l)]$$

It follows that when $\theta_s < \theta_b$, we are unable to attain the CIE surplus.

Now suppose that $p^* = p^l$ and $p_* = p^h$. The buyer can unilaterally raise the price in state h from p^h to p^l. The following SPIC constraint prevents him from doing so:

$$h - p^h - t_s(p^h) \geq h - p^l - t_s(p^l)$$

Therefore, $p^h + t_s(p^h) - p^l - t_s(p^l) \leq 0$. It follows that when $\theta_s > \theta_b$, we are unable to attain the CIE surplus.

It can be shown that the CIE surplus can be attained if the assignment of trading prices to states depends on the identity of the agent with the highest θ. This means that CIE bets cannot be purely speculative. It turns out that although we are unable to apply our main result, the same methods can be adapted to demonstrate that the CIE surplus is implementable for every F. The key to this adaptation is to view the trading price p^ω as part of the transfer that takes place in state ω (and as such to allow it to depend on $\boldsymbol{\theta}$), and then use the SPIC constraints to derive bounds on $p^h + t_s(p^h) - p^l - t_s(p^l)$, rather than on $t_s(p^h) - t_s(p^l)$. For the sake of brevity, we omit the proof of this claim.

Multilateral speculation problems

We have restricted attention to bilateral speculation problems. Extending the model to games with more than two agents is straightforward. However, Proposition 2 ceases to hold in this case. For instance, suppose that the partition X is the finest possible - that is, the agents can sign bets that condition on the second-period action profile. Then, under mild assumptions on the bare-game payoff structure, infinite bets become possible, by letting agents 1 and 2 bet on agent 3's action. Agents 1 and 2 are thus unable to manipulate the bet's outcome, and therefore the stakes of their bet are unlimited. The only problem is to provide agent 3 with incentives to play different actions in the two states. But since agents 1 and 2 earn unlimited speculative gains, they can use these gains to provide the necessary incentives.

Our approach, however, remains fruitful in some special cases. One simple case is when the partition X consists of only two cells. For instance, suppose that in period 2, the agents play a voting game. There are two candidates, A and B. Unless all agents vote for candidate B, the elected candidate is A. The agents have quasi-linear, state-dependent utility. In state u, $u_i(A) \geq u_i(B)$ for all i. In state v, $v_i(B) \geq v_i(A)$ for all i. In period 1, the agents can only bet on the identity of the elected candidate. The CIE surplus can be attained only if all agents vote B in state v, and at least two agents vote A in state u.

The structure of CIE bets is such that agent $i^* = \arg\min_i \theta_i$ - i.e., the agent who has the biggest faith in the election of candidate B - essentially signs a bilateral side bet with every other agent. The stakes of the bilateral bet between i^* and j are $v_j(B) - v_j(A)$, namely j's cost of unilaterally imposing A as the elected candidate in state v. It can be shown that the problem of implementing the CIE surplus in this case is equivalent to the problem of implementing efficient dissolution of an n-player partnership of size $\Sigma_i[v_i(B) - v_i(A)]$, in which the share of partner i in the jointly owned asset is $v_i(B) - v_i(A)$. Thus, using Propositions 1-3 in CGK, it can be shown that as the utility differences $v_i(B) - v_i(A)$ become more symmetric across agents, it becomes possible to implement the CIE surplus for a larger set of distributions from which the agents' priors are drawn.

Speculation problems with more than two states

Our model of bilateral speculation problems assumes two states of Nature. This is a greatly simplifying device, since it implies that an agent's type is a scalar. When we extend the model to environments with $K > 2$ states of Nature, an agent's type is an element in the K-dimensional simplex, and therefore the problem of implementing CIE bets is a mechanism-design problem with multi-dimensional types.

The idea that CIE bets may be formally equivalent to efficient dissolution of a partnership may be extended to these environments. However, new considerations arise. First, the partnership may involve up to $K-1$ assets, and the parties' ownership shares may be asset-specific. Second, the values that a party attaches to any pair of these assets are negatively correlated. Third, the bilateral speculation problem may be characterized by a large number of SPIC constraints, which translate into additional constraints on the final allocations of the assets in the analogous multi-asset partnership dissolution problem (for instance, giving different parties full ownership of different assets may be infeasible).

Thus, when there are more than two states, our model may be formally equivalent to a multi-asset partnership dissolution problem, with constraints on the agents' valuations and the set of feasible final allocations. A general characterization of this equivalence lies beyond the scope of the present paper.

Non-common priors versus state-dependent utility

Our main result utilizes a formal equivalence between our model of speculative trade and a model of trade motivated by differences in tastes. The question arises, whether our model could be re-interpreted as a standard model in the first place, since it is well-known that state-dependent utility and subjective probability are impossible to distinguish behaviorally. At first glance, the answer is affirmative: our model is behaviorally equivalent to a model in which every agent i assigns probability $\frac{1}{2}$ to each state, and his utility function is multiplied by a state-dependent constant (θ_i in one state and $1-\theta_i$ in the other state). However, this re-interpretation requires us to make two assumptions: *(i)* the agents' utility from money is *state-dependent*; *(ii)* the agents' trade-off between money and bare-game outcomes is *state-independent*. We find it extremely hard to imagine a reasonable justification for such preferences. Therefore, θ_i is more convincingly interpreted as a prior belief than as a taste parameter.

Related literature

This paper follows up Eliaz and Spiegler (2005,2006), in which we analyze the problem of designing a profit-maximizing menu of contracts for a monopolist facing a population of consumers who differ in their ability to forecast their future tastes. In Eliaz and Spiegler (2006), 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 are necessary for price discrimination.

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 payments that modify the second-period game. We are aware of a number of precedents for this aspect of our paper. Allaz and Vila (1993) 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. Jackson and Wilkie (2005) study two-stage games, in which players commit to unilateral transfers conditional on the outcome of a later "bare game". They study the properties of subgame perfect equilibria in such games. Both works assume away any uncertainty regarding second-period payoffs.

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.

The partnership dissolution model studied by CGK was taken up by Fieseler, Kittsteiner and Moldovanu (2003) and Jehiel and Pauzner (2004), who extended the informational structure to allow for interdependent valuations. Neeman (1999) studies the closely related problem of characterizing the structure of property rights for which voluntary bargaining can resolve a public good problem efficiently.

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

[5] Eliaz, Kfir and Ran Spiegler (2006): "Contracting with Diversely Naive Agents," *Review of Economic Studies* **73**, 689-714.

[6] Fieseler, Karsten, Thomas Kittsteiner and Benny Moldovanu (2003): "Partnerships, Lemons and Efficient Trade," *Journal of Economic Theory*, **113**, 223-234.

[7] Jackson, Matthew and Simon Wilkie (2005): "Endogenous Games and Mechanisms: Side Payments among Players," *Review of Economic Studies*, **72**, 543-566.

[8] Jehiel, Philippe and Ady Pauzner (2004): "Partnership Dissolution with Interdependent Values", *Rand Journal of Economics*, forthcoming.

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

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

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

[12] Neeman, Zvika (1999): "Property Rights and Efficiency of Voluntary Bargaining under Asymmetric Information," *Review of Economic Studies*, **66**, 679-691.

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

[14] Yildiz, Muhamet (2004): "Wishful Thinking in Strategic Environments," *Review of Economic Studies*, forthcoming.

Appendix: Proofs

Proof of Proposition 2

Consider some CIE tuple (a^u, a^v, t). By the SPIC constraints and budget-balancedness,

$$\begin{aligned}
t_1[x(a_1^v, a_2^u)] - t_1[x(a_1^u, a_2^u)] &\leq u_1(a_1^u, a_2^u) - u_1(a_1^v, a_2^u) \\
t_1[x(a_1^u, a_2^u)] - t_1[x(a_1^u, a_2^v)] &\leq u_2(a_1^u, a_2^u) - u_2(a_1^u, a_2^v) \\
t_1[x(a_1^v, a_2^v)] - t_1[x(a_1^v, a_2^u)] &\leq v_2(a_1^v, a_2^v) - v_2(a_1^v, a_2^u) \\
t_1[x(a_1^u, a_2^v)] - t_1[x(a_1^v, a_2^v)] &\leq v_1(a_1^v, a_2^v) - v_1(a_1^u, a_2^v)
\end{aligned}$$

These inequalities together imply:

$$\begin{aligned} t_1[x(a_1^v, a_2^v)] - t_1[x(a_1^u, a_2^u)] &\leq [u_1(a_1^u, a_2^u) - u_1(a_1^v, a_2^u)] + [v_2(a_1^v, a_2^v) - v_2(a_1^v, a_2^u)] \\ t_1[x(a_1^u, a_2^u)] - t_1[x(a_1^v, a_2^v)] &\leq [u_2(a_1^u, a_2^u) - u_2(a_1^u, a_2^v)] + [v_1(a_1^v, a_2^v) - v_1(a_1^u, a_2^v)] \end{aligned}$$

Because u and v are bounded, $\theta_i u_i(a^u) + (1-\theta_i)v_i(a^v)$ is finite for each agent i. Moreover, the R.H.S in the last two inequalities are finite. But this means that $t_1[x(a_1^u, a_2^u)] - t_1[x(a_1^v, a_2^v)]$ is finite. ■

Proof of Proposition 3

Assume that the CIE surplus is attained by pure speculation and consider some purely speculative CIE tuple $(a^u, a^v, t'(\boldsymbol{\theta}))$. Then for all $\boldsymbol{\theta} = (\theta_1, \theta_2)$, the bet $t'(\boldsymbol{\theta})$ maximizes

$$(\theta_1 - \theta_2) \cdot [t_1'(x^u) - t_1'(x^v)] \tag{4}$$

subject to the SPIC constraints.

We proceed in two steps. First, we show that we can construct a bet t that satisfies (3) as well as the SPIC constraints. Second, we show that (3) is necessary for maximizing (4) subject to the SPIC constraints.

The first step of our proof relies on the following lemma.

Lemma 1 *Let a^u and a^v be pure-strategy NE of $G(u)$ and $G(v)$, respectively, and let t be a bet that satisfies*

$$t_1(y) - t_1(x^v) = \min[d_1(a^v \to y), \hat{D} + d_1(a^u \to y)] \tag{5}$$

for all $y \in X$, where $\hat{D} \in \{D_2(a^u, a^v), -D_1(a^u, a^v)\}$. Then a^u and a^v are also pure-strategy NE of $G(u,t)$ and $G(v,t)$, respectively.

Proof of Lemma 1. The SPIC constraints, which ensure that a^u and a^v are also pure-strategy NE of $G(u,t)$ and $G(v,t)$, may be summarized by the following inequalities (which use budget-balancedness). For every $y \in X$:

$$\begin{aligned} t_1(y) - t_1(x^u) &\leq d_1(a^u \to y) && (6) \\ t_1(y) - t_1(x^v) &\leq d_1(a^v \to y) && (7) \\ t_1(x^v) - t_1(y) &\leq d_2(a^v \to y) && (8) \\ t_1(x^u) - t_1(y) &\leq d_2(a^u \to y) && (9) \end{aligned}$$

Suppose that $d_1(a^u \to y) + \hat{D} \leq d_1(a^v \to y)$. Then, by (5), inequalities (6) and (7) are satisfied. Assume that (8) is violated. Then, by (5):

$$-\hat{D} > d_1(a^u \to y) + d_2(a^v \to y) \tag{10}$$

If $\hat{D} = -D_1(a^u, a^v)$, then by the definition of $D_1(a^u, a^v)$, the L.H.S of (10) cannot exceed its R.H.S., a contradiction. If $\hat{D} = D_2(a^u, a^v)$, then by our assumption that $[a^u, a^v, t'(\boldsymbol{\theta})]$ is a purely speculative CIE tuple,

$$-D_2(a^u, a^v) \leq 0 \leq d_1(a^u \to y) + d_2(a^v \to y)$$

contradicting (10). Therefore, (8) must hold. Finally, to see that (9) is satisfied, note that the L.H.S of this inequality is equal to $-d_1(a^u \to y)$ and by our pure speculation assumption,

$$-d_1(a^u \to y) \leq 0 \leq d_2(a^u \to y)$$

Alternatively, suppose that $d_1(a^u \to y) + \hat{D} > d_1(a^v \to y)$. Then, by (5), inequalities (6) and (7) are satisfied. Assume that (9) is violated. Then, by (5):

$$\hat{D} > d_1(a^v \to y) + d_2(a^u \to y) \tag{11}$$

If $\hat{D} = D_2(a^u, a^v)$, then by definition, it cannot exceed the R.H.S. of (11), a contradiction. If $\hat{D} = -D_1(a^u, a^v)$, then by our assumption that $[a^u, a^v, t'(\boldsymbol{\theta})]$ is a purely speculative CIE tuple,

$$-D_1(a^u, a^v) \leq 0 \leq d_1(a^v \to y) + d_2(a^u \to y)$$

contradicting (11). Therefore, (9) must hold. Finally, (8) follows from our pure speculation assumption, which implies that

$$-d_1(a^v \to y) \leq 0 \leq d_2(a^v \to y)$$

This concludes the proof of the lemma. □

Construct a bet t that satisfies (5) for every $y \in X$. Note that the only restriction on $\hat{D}$ is that it has only two possible values, $D_2(a^u, a^v)$ or $-D_1(a^u, a^v)$. Let $\hat{D} = D^*(a^u, a^v \mid \boldsymbol{\theta})$. Then for $y = x^u$, the bet t satisfies (3). By Lemma 1, t also satisfies the SPIC constraints. This completes the first step of our proof.

Our next step is to show that if t satisfies the SPIC constraints, then:

$$-D_1(a^u, a^v) \le t_1(x^u) - t_1(x^v) \le D_2(a^u, a^v) \tag{12}$$

The SPIC constraints, summarized by (6)-(9), imply that for every $y \in X$:

$$-d_1(a^u \to y) - d_2(a^v \to y) \le t_1(x^u) - t_1(x^v) \le d_2(a^u \to y) + d_1(a^v \to y)$$

But this boils down to (12). Therefore, (3) is necessary for constrained interim-efficiency. ■

Proof of Proposition 4

We proceed in two steps. First, let us show that implementation of the CIE surplus is **sufficient** for efficient dissolution of the partnership $\langle D_1(a^u, a^v), D_2(a^u, a^v), F \rangle$. Assume the CIE surplus of $\langle (u, v), G, X, F \rangle$ is implementable .Consider the following mechanism: for $i = 1, 2$, and for every pair of reports $\hat{\boldsymbol{\theta}}$,

$$\begin{aligned} q_i(\hat{\boldsymbol{\theta}}) &= D_i(a^u, a^v) + t_i^u(\hat{\boldsymbol{\theta}}) - t_i^v(\hat{\boldsymbol{\theta}}) \\ m_i(\hat{\boldsymbol{\theta}}) &= t_i^v(\hat{\boldsymbol{\theta}}) \end{aligned}$$

where, for notational ease, we let $t_i^\omega(\hat{\boldsymbol{\theta}}) \equiv t_i(x^\omega | \, \hat{\boldsymbol{\theta}})$ for $\omega = u, v$. Because $t_i^u(\hat{\boldsymbol{\theta}})$ and $t_i^v(\hat{\boldsymbol{\theta}})$ satisfy (EFF), (PS-SPIC), (IC) and (IR) it follows that the mechanism $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ has the following properties. First, by (EFF), whenever $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$,

$$q_1(\boldsymbol{\theta}) = \begin{cases} D_1(a^u, a^v) + D_2(a^u, a^v) & if \quad \theta_1 \ge \theta_2 \\ 0 & if \quad \theta_1 < \theta_2 \end{cases}$$

Hence, $q(\hat{\boldsymbol{\theta}})$ satisfies (EFF*). Second, by (IC) and (IR), we have that for $i = 1, 2$, and $\theta_i' \in [0, 1]$,

$$\theta_i[Q_i(\theta_i) - D_i(a^u, a^v)] + M_i(\theta_i) \ge \theta_i[Q_i(\theta_i') - D_i(a^u, a^v)] + M_i(\theta_i')$$

and

$$\theta_i[Q_i(\theta_i) - D_i(a^u, a^v)] + M_i(\theta_i) \ge 0$$

These two inequalities imply that $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ satisfies (IC*) and (IR*).

We now show that implementation of the CIE surplus is **necessary** for efficient dissolution of the partnership $\langle D_1(a^u, a^v), D_2(a^u, a^v), F \rangle$. Let $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ be a direct mechanism that efficiently dissolves the partnership $\langle D_1(a^u, a^v), D_2(a^u, a^v), F\rangle$. Then, for every realization of $\boldsymbol{\theta} \in [0,1]^2$, this mechanism satisfies (EFF*), (IC*) and (IR*). Now consider a bilateral speculation problem $\langle (u,v), G, X, F\rangle$ where the CIE surplus is attained by pure speculation and sustained by (a^u, a^v, t). By the proof of Proposition 3, t satisfies (5), without loss of generality.

Let $t(x \mid \hat{\boldsymbol{\theta}})$ be a direct mechanism for $\langle (u,v), G, X, F\rangle$ such that for every $i = 1, 2$, and for all profiles of reports $\hat{\boldsymbol{\theta}}$:

$$t_1(x^v \mid \hat{\boldsymbol{\theta}}) = m_1(\hat{\boldsymbol{\theta}}) \tag{13}$$

and for every $y \neq x^v$:

$$t_1(y \mid \hat{\boldsymbol{\theta}}) - t_1(x^v \mid \hat{\boldsymbol{\theta}}) = \min[d_1(a^v \to y), d_1(a^u \to y) + q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v)] \tag{14}$$

Because $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ satisfies (EFF*), $q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v) = D^*(a^u, a^v \mid \boldsymbol{\theta})$. In particular, this means that $q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v)$ is equal to either $D_2(a^u, a^v)$ or $-D_1(a^u, a^v)$. In either case, if $y = x^u$, then by the definition of $D_1(a^u, a^v)$ and $D_2(a^u, a^v)$, $d_1(a^u \to x^u) + q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v)$ cannot exceed $d_1(a^v \to x^u)$. Hence,

$$t_1(x^u \mid \hat{\boldsymbol{\theta}}) - t_1(x^v \mid \hat{\boldsymbol{\theta}}) = q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v) \tag{15}$$

The observation that $q_1(\hat{\boldsymbol{\theta}}) - D_1(a^u, a^v) = D^*(a^u, a^v \mid \boldsymbol{\theta})$ implies that equation (14) becomes equation (3). This means that when $y = x^u$, $t(x \mid \hat{\boldsymbol{\theta}})$ satisfies (EFF). By Lemma 1, this also means that $t(x \mid \hat{\boldsymbol{\theta}})$ satisfies (PS-SPIC). It remains to show that $t(x \mid \hat{\boldsymbol{\theta}})$ satisfies (IC) and (IR). Since $(q(\hat{\boldsymbol{\theta}}), m(\hat{\boldsymbol{\theta}}))$ satisfies (IC*) and (IR*), the following inequalities must hold for $i = 1, 2$, and for all $\theta_i' \in [0, 1]$,

$$\begin{aligned}
\theta_i Q_i(\theta_i) + M_i(\theta_i) &\geq \theta_i Q_i(\theta_i') + M_i(\theta_i') \\
\theta_i Q_i(\theta_i) + M_i(\theta_i) &\geq \theta_i D_i(a^u, a^v)
\end{aligned}$$

Rewriting these inequalities, we obtain

$$\begin{aligned}
\theta_i[Q_i(\theta_i) - D_i(a^u, a^v)] + M_i(\theta_i) &\geq \theta_i[Q_i(\theta_i') - D_i(a^u, a^v)] + M_i(\theta_i') \\
\theta_i[Q_i(\theta_i) - D_i(a^u, a^v)] + M_i(\theta_i) &\geq 0
\end{aligned}$$

By the definitions of $Q_i(\theta_i')$ and $M_i(\theta_i')$, and the relation between $t(x \mid \hat{\boldsymbol{\theta}})$ and $q_1(\hat{\boldsymbol{\theta}})$ given by (15), the last two inequalities imply (IC) and (IR), respectively. ■

Proof of Proposition 5

We prove the result stepwise.

Step 1. For every t, it is impossible to sustain a market price $p^\omega > c^\omega$ in a NE of $G(\omega, t)$.

Proof. Let (a_1^ω, a_2^ω) be a NE of $G(\omega, t)$ that satisfies $\min\{a_1^\omega, a_2^\omega\} = p^\omega > c^\omega$. Then, for all i and for all $\varepsilon > 0$,

$$s_i(a_1^\omega, a_2^\omega) \cdot (p^\omega - c^\omega) + t_i(p^\omega) \geq p^\omega - \varepsilon - c^\omega + t_i(p^\omega - \varepsilon)$$

where

$$s_i(a_1^\omega, a_2^\omega) = \begin{cases} 1 & if \quad a_i < a_j \\ \frac{1}{2} & if \quad a_i = a_j \\ 0 & if \quad a_i > a_j \end{cases}$$

Summing over i and using budget-balancedness, we obtain:

$$p^\omega - c^\omega \geq 2(p^\omega - \varepsilon - c^\omega)$$

for all $\varepsilon > 0$. But this implies that (p^ω, p^ω) is a NE of $G(\omega)$, a contradiction.

Step 2. If the CIE surplus is attained by pure speculation, then the CIE surplus is

$$|\theta_1 - \theta_2| \cdot (c^H - c^L) \tag{16}$$

Proof. If the CIE surplus is attained by pure speculation, then any CIE tuple (a^L, a^H, t) satisfies $p^\omega = c^\omega$. Since this means that the sellers' bare-game payoff is zero in both states, the CIE surplus may be written as $(\theta_1 - \theta_2)(t_1^L - t_1^H)$, where $t_1^L \equiv t_1(c^L)$ and $t_1^H \equiv t_1(c^H)$. In state H, each seller can unilaterally lower the price to c^L. This deviation is not profitable if the following SPIC constraint holds: for every seller i, $t_i^H \geq c^L - c^H + t_i^L$. By budget-balancedness,

$$c^L - c^H \leq t_1^L - t_1^H \leq c^H - c^L \tag{17}$$

Therefore, the CIE surplus is bounded from above by $|\theta_1 - \theta_2| \cdot (c^H - c^L)$. To see that this expression can be attained, define $t(x \mid \boldsymbol{\theta})$ as follows. When $\theta_1 \geq \theta_2$, let $t_1^L - t_1^H = c^H - c^L$. Conversely, when $\theta_1 < \theta_2$, let $t_1^L - t_1^H = c^L - c^H$. In both cases, let $t_1(p) = t_1^L$ for every $p \leq c^L$, and let $t_1(p) = t_1^H$ for every $p > c^L$. Because t is a step function, and because a^ω is a NE in $G(\omega)$, all SPIC constraints hold, and the surplus is $|\theta_1 - \theta_2| \cdot (c^H - c^L)$.

Step 3. Total interim surplus evaluated at any (a^L, a^H, t), with t satisfying the SPIC, is at most $|\theta_1 - \theta_2| \cdot (c^H - c^L)$.

Proof. Denote $p^* \equiv \max\{p^H, p^L\}$ and $p_* \equiv \min\{p^H, p^L\}$. Let ω^* and ω_* denote the states in which p^* and p_* occur, and let c^* and c_* denote the marginal costs in states ω^* and ω_* respectively. Let θ_i^* be seller i's prior on ω^*, and denote his market share in ω^* by s_i.

By Step 1, $p^* \leq c^*$ and $p_* \leq c_*$. If $p^* = c^*$ and $p_* = c_*$, then by Step 2, the proof is complete. Now assume that one of these inequalities holds strictly. Because both sellers can unilaterally lower the market price from p^* to p_* in state ω^*, the following SPIC constraints must hold:

$$\begin{aligned} s_1 \cdot (p^* - c^*) + t_1(p^*) &\geq p_* - c^* + t_1(p_*) \\ s_2 \cdot (p^* - c^*) + t_2(p^*) &\geq p_* - c^* + t_2(p_*) \end{aligned}$$

Using budget-balancedness, we obtain:

$$p_* - c^* - s_1 \cdot (p^* - c^*) \leq t_1(p^*) - t_1(p_*) \leq s_2 \cdot (p^* - c^*) + c^* - p_* \tag{18}$$

Suppose $p_* = c_*$. Then, $p^* < c^*$ and $c_* < c^*$. Because $c^H > c^L$ it follows that $\omega^* = H$. But in this case the SPIC constraints given by (18) imply that total surplus is less than (16). It follows that $p_* < c_*$. This means that there is exactly one seller i who plays $a_i = p_*$ in state ω_*. If both sellers played p_*, then either one of them could deviate upward. This deviation would leave market price (and therefore the transfers) unaffected, but it would save the deviator a bare-game loss. Without loss of generality, assume that seller 1 sustains the market price p_* in state ω_*. Let $a_2 > p_*$ denote seller 2's action in this state.

It follows that the sellers' total interim surplus is given by the following expression:

$$(p^* - c^*) \cdot (s_1\theta_1^* + s_2\theta_2^*) + (1 - \theta_1^*) \cdot (p_* - c_*) + (\theta_1^* - \theta_2^*) \cdot [t_1(p^*) - t_1(p_*)]$$

Note that the first two terms are non-positive, and one of them is strictly negative, by assumption. Therefore, if we prove that the third term does not exceed (16), we complete the proof.

Suppose that $\theta_2^* \geq \theta_1^*$. Then, by (18), total interim surplus is bounded from above by

$$(p^* - c^*) \cdot (s_1\theta_1^* + s_2\theta_2^*) + (1 - \theta_1^*) \cdot (p_* - c_*) + (\theta_1^* - \theta_2^*) \cdot [p_* - c^* - s_1 \cdot (p^* - c^*)]$$

Because $s_2 = 1 - s_1$, this expression may be rewritten as

$$(p^* - c^*) \cdot \theta_2^* + p_* \cdot (1 - \theta_2^*) - c_* \cdot (1 - \theta_1^*) + c^* \cdot (\theta_2^* - \theta_1^*)$$

Since $c^* \leq c^H$, this expression is at most

$$(p^* - c^*) \cdot \theta_2^* + p_* \cdot (1 - \theta_2^*) - c_* \cdot (1 - \theta_1^*) + c^H \cdot (\theta_2^* - \theta_1^*)$$

By adding and subtracting $c_*\theta_2^*$, we may rewrite this expression as

$$(p^* - c^*) \cdot \theta_2^* + (p_* - c_*) \cdot (1 - \theta_2^*) + (\theta_2^* - \theta_1^*) \cdot (c^H - c_*)$$

Because $p^* \leq c^*$ and $p_* < c_*$, the above expression is strictly below $(\theta_2^* - \theta_1^*) \cdot (c^H - c_*)$. But since $\theta_2^* \geq \theta_1^*$ and $c_* \geq c^L$,

$$(\theta_2^* - \theta_1^*) \cdot (c^H - c_*) < (\theta_2^* - \theta_1^*) \cdot (c^H - c^L)$$

Our assumption that $\theta_2^* \geq \theta_1^*$ implies that whether $\omega^* = \omega^H$ or $\omega^* = \omega^L$, the R.H.S. of the above inequality is $|\theta_1 - \theta_2| \cdot (c^H - c^L)$.

Now suppose that $\theta_1^* > \theta_2^*$. In addition to the SPIC constraints given by (18), there is an additional SPIC constraint, which prevents seller 1 from raising the market price from p_* to a_2. There are three cases to consider.

Case 1: $a_2 < p^*$. Seller 1 can deviate from $a_1 = p_*$ to $a_1' \in (a_2, p^*)$. The SPIC constraint that prevents him from doing so is

$$p_* - c_* + t_1(p_*) \geq t_1(a_2)$$

But note that in state ω^*, seller 2 can unilaterally lower the market price from p^* to a_2. The SPIC constraint that prevents him from doing so is

$$s_2 \cdot (p^* - c^*) - t_1(p^*) \geq a_2 - c^* - t_1(a_2)$$

Combining these two constraints, we obtain

$$t_1(p^*) - t_1(p_*) \leq c^* - c_* + p_* - a_2 + s_2 \cdot (p^* - c^*)$$

but the R.H.S of this inequality is lower than $c^H - c^L$.

Case 2: $a_2 > p^*$. Seller 1 can deviate from $a_1 = p_*$ to $a_1' = p^*$. The SPIC constraint that prevents him from doing so is

$$p_* - c_* + t_1(p_*) \geq p^* - c_* + t_1(p^*)$$

This constraint implies $t_1(p^*) - t_1(p_*) \leq p_* - p^* < 0 < c^H - c^L$.

Case 3: $a_2 = p^*$. Seller 1 can deviate from $a_1 = p_*$ to $a_1' > a_2$ or $a_1' = a_2$. The SPIC constraint that prevents him from carrying out either of these deviations is

$$p_* - c_* + t_1(p_*) \geq \max[0, \frac{1}{2}(p^* - c_*)] + t_1(p^*)$$

This constraint implies $t_1(p^*) - t_1(p_*) \leq p_* - c_* - \max[0, \frac{1}{2}(p^* - c_*)] < 0 < c^H - c^L$.

We have thus established that the SPIC constraints that result from setting $p_* < c_*$ imply that total interim surplus is below (16). ■

Proof of Proposition 6

If there is no trade in both states, there is no scope for speculation. Let ω denote a state with trade and let p^ω denote the market price in this state. Each agent can unilaterally impose no trade in ω (the seller can submit an ask price above the buyer's bid price, and the buyer can submit a bid price below the seller's ask price). Therefore, the SPIC constraints that prevent these deviations are:

$$\begin{aligned} p^\omega - c + t_s(p^\omega) &\geq t_s(NT) \\ \omega - p^\omega - t_s(p^\omega) &\geq -t_s(NT) \end{aligned} \tag{19}$$

Hence,

$$c - p^\omega \leq t_s(p^\omega) - t_s(NT) \leq \omega - p^\omega \tag{20}$$

Suppose there is trade in only one state. Then the total interim-expected surplus is given by

$$\pi_s(p^\omega - c) + \pi_b(\omega - p^\omega) + (\pi_s - \pi_b) \cdot [t_s(p^\omega) - t_s(NT)] \tag{21}$$

where π_i denotes agent i's prior on ω.

By (20), total surplus cannot exceed $\max\{\pi_b, \pi_s\} \cdot (\omega - c)$. Since we are free to choose the state in which trade occurs, we can set $\omega = h$. Therefore, total surplus is at most:

$$\max\{\pi_b, \pi_s\} \cdot (h - c) \tag{22}$$

Now suppose that trade occurs in both states. Then the inequalities in (19) are the SPIC constraints that prevent each agent from unilaterally imposing no trade in each state. Hence,

$$p^l - p^h + c - l \leq t_s(p^h) - t_s(p^l) \leq p^l - p^h + h - c \tag{23}$$

while total surplus is

$$-c + \theta_b h + (1 - \theta_b)l + (\theta_s - \theta_b) \cdot [p^h - p^l + t(p^h) - t(p^l)]$$

By (23), total surplus is bounded from above by:

$$\max\{\theta_b, \theta_s\}(h - c) + (1 - \min\{\theta_b, \theta_s\})(l - c) \tag{24}$$

which is below (22), since $l < c$. It follows that the value of the CIE surplus is given by (22), and that the market outcome induced by the CIE surplus is ex-post efficient.

We now proceed to prove part *(ii)*. It is easy to see that when we plug the values of p^h and $t_s(T) - t_s(NT)$, as stated in part *(ii)*, into expression (2), we obtain (22). Therefore, it remains to show that the SPIC constraints are satisfied. Consider state h. The buyer's payoff is $\frac{h-c}{2} - t_s(T)$. If the buyer raises his bid price, he raises the market price, and therefore loses in terms of bare-game payoffs, without affecting the transfer. If he lowers his bid price, he imposes no trade, in which case his net payoff is $-t_s(NT)$. Because $t_s(T) - t_s(NT)$ is $\frac{h-c}{2}$ if $\theta_s > \theta_b$ and $\frac{c-h}{2}$ otherwise, this deviation is not profitable. The seller's payoff is $\frac{h-c}{2} + t_s(T)$. If he lowers his ask price, he loses in terms of bare-game payoffs, without affecting the transfer. If he raises his ask price, he imposes no trade, in which case his net payoff is $t_s(NT)$. It follows that neither deviation is profitable. Now consider state l. Since p_s^l is arbitrarily high and p_b^l is arbitrarily low, neither agent has an incentive to enforce trade unilaterally. ■