



Bargaining over bets

Kfir Eliaz^{a,*}, Ran Spiegler^b

^a Department of Economics, Brown University, Box B, Providence, RI 02912, USA

^b University College London, Department of Economics, Gower Street, London WC1E 6BT, UK

Received 14 June 2007

Available online 1 July 2008

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.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C79; D84; D86

Keywords: Speculative trade; Mechanism design; Bets; Pre-game contracts; Non-common priors; Optimism

1. Introduction

In many situations people hold different opinions about how the future will unfold. While these differences may be partially explained by asymmetric information, some differences may still persist even if the individuals were to share all their information. This could be the result of inherent biases that people have in forming their beliefs, such as optimism and pessimism, or overconfidence in one's ability to process information. For example, a die-hard Mets fan, who watches or reads the same sports commentary as a die-hard Yankees fan, would still disagree with the latter on the likelihood that his team would win the World Series. Another example includes entrepreneurs, who are often inherently more optimistic than venture capitalists with regards to the profitability of their start-ups. Similarly, a group of entrepreneurs starting a business often have different beliefs about the likelihood of future success. A fashion retailer may believe that his daily contact with consumers makes him better capable of predicting future demand than his supplier, who holds the opposite view owing to his experience with a large number of retailers. Finally, traders in a market often arrive at different conclusions about what the current economic data implies with regards to future prices.

* Corresponding author.

E-mail addresses: Kfir_Eliaz@Brown.edu (K. Eliaz), r.spiegler@ucl.ac.uk (R. Spiegler).

URLs: http://www.econ.brown.edu/fac/Kfir_Eliaz/index.htm (K. Eliaz), <http://www.homepages.ucl.ac.uk/~uctprsp> (R. Spiegler).

In situations such as these, the parties involved can make speculative gains by betting on future events. These bets may take an explicit form such as bets made on sporting events. They could be somewhat less explicit as in form of financial derivatives. But they could also be implicit in the form of contracts that the parties may write. For example, an optimistic entrepreneur would be willing to exchange cash flow rights when performance is low against claims when performance is high (see Landier and Thesmar, in press); an optimistic partner in a business venture would prefer to lower the salary he withdraws in return for more stock options, while a pessimistic partner would hold the opposite preference; and an optimistic retailer would agree to commit in advance to a large stock for a relatively low price and pay a high penalty for any remaining units that are not ordered (see Bazerman and Gillespie, 1999). These observations suggest that institutions, which allow parties with heterogeneous beliefs to bet on the future, are quite prevalent. To better understand these institutions, we need first to understand what are the maximal speculative gains that can be made in the relevant situations, and whether these gains can be attained when the parties' beliefs are not observed.

This paper addresses these questions with a simple two-period model, which we call a “Bilateral Speculation Problem.” 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” (CIE) bet and ask whether it can be implemented in Bayesian equilibrium by some mechanism.

Our model generalizes the framework of Eliaz and Spiegel (2007), which took the first step towards addressing the above questions in a context of a simple example. The example is a special case of our model in which only a single agent takes an action and has only two possible actions to choose from. Eliaz and Spiegel (2007) show that in this example, implementing a CIE bet between two parties is equivalent to efficiently allocating an asset, which each party initially owns some share of. This latter problem, first studied by Cramton et al. (1987), extends the buyer–seller allocation problem of Myerson and Satterthwaite (1983) to more general ownership structures, namely “partnerships” (hence, it is also termed “efficient dissolution of a partnership”). A manipulable bet may be viewed as an asset—an entitlement to receive a prize conditional on a random event—whose value and initial ownership structure are determined by the underlying manipulation costs.

Our main result in this paper provides the conditions under which the equivalence to the partnership dissolution problem extends to the more general framework studied here. More precisely, we show that the equivalence holds whenever the CIE surplus can be attained by a “purely speculative” bet, i.e., a bet that induces second-period behavior, which is independent of the agents' priors, and which may be sustained in the absence of bets (formally, the bet induces the agents to play a Nash equilibrium of the original game in the second period). We demonstrate the usefulness of this equivalence result 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.

A key feature of the environment we study is that parties bet on outcomes that at least one of them can manipulate. This feature is common to (almost) all of the above examples: an entrepreneur can manipulate the profitability of his business venture by insisting on pursuing his initial plans despite new information that may call for a change; a partner in a business can manipulate the firm's performance by trying excessively risky strategies; a retailer can manipulate the amount he pays to his supplier by ordering quantity in excess of real demand; and a large trader can manipulate future prices by submitting large buy or sell orders on the spot market (see Newbery, 1984). Finally, a sporting bet can be manipulated if one of the parties actually competes in the event.

Our assumption that bets are manipulable, means that agents would not agree to bet on arbitrarily large amounts. 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, as mentioned above, there are many real-life situations in which agents with heterogeneous beliefs bet on outcomes they can manipulate. 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 paper is organized as follows. The next section presents our model and our main result. In Section 3 we consider two simple market applications of our result. Section 4, which concludes the paper, discusses various aspects of our model and also reviews some of the related literature.

2. 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 . We refer to G as the “bare game.” 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 $u_i : A_1 \times A_2 \rightarrow \mathbb{R}$, while in the second state this function is $v_i : A_1 \times A_2 \rightarrow \mathbb{R}$. Let $G(\omega)$ denote the second-period game played in state $\omega \in \{u, v\}$. We assume that $G(\omega)$ has a pure-strategy Nash equilibrium 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 $G(\omega, t)$.

We illustrate the model with the following simple example. Consider two team members, who each needs to decide whether to exert effort (action h) or not (action l) on some joint project. Hence, G is a 2×2 game in which $A_1 = A_2 = \{h, l\}$. Unless both exert effort, the project is not completed. In state u there are two pure-strategy Nash Equilibria, one in which both agents exert no effort and another, Pareto superior equilibrium, where both exert effort (i.e., $u_i(h, h) > u_i(l, l)$ for $i = 1, 2$). In state v , it is a dominant strategy for each agent to exert no effort (i.e., $v_i(l, l) > v_i(h, l)$ and $v_i(l, h) > v_i(h, h)$ for $i = 1, 2$). Each member's decision and payoff are observed only by him. Hence, period 1 bets can be made contingent only on whether or not the project is completed, i.e., $X = \{x^h, x^l\}$ where $x^h = \{(h, h)\}$ and $x^l = \{(l, l), (h, l), (l, h)\}$. These bets may be interpreted as profit-sharing agreements: the two agents would like to redistribute the project's profits (or losses) among themselves, but since their payoffs are not verifiable they can only specify budget balanced transfers as a function of whether or not the project is completed.

Given a bilateral speculation problem, such as the one described above, we ask the following questions. First, for a given pair of prior beliefs, what are the largest gains from bets that the two agents can make? Second, can these gains be attained when prior beliefs are private information?

Let us begin with our first question. Consider a bet t , and suppose that both agents expect the action profiles in states u and v to be $a^u = (a_1^u, a_2^u)$ and $a^v = (a_1^v, a_2^v)$ respectively. Then agent i 's interim expected payoff from (a^u, a^v, t) is

$$\theta_i[u_i(a^u) + t_i(x(a^u))] + (1 - \theta_i)[v_i(a^v) + t_i(x(a^v))]$$

We use the term “interim” to highlight the analogy with standard models of trade in which an agent's type is a preference parameter. As in the standard models, the agent's expected payoff is calculated after he learns his type which is drawn from a common distribution F . In another sense, the term is inappropriate because the agent's type is a prior belief and F can be viewed as the agent's second-order belief of the opponent's prior.

The sum of the parties' interim expected payoffs can be conveniently written as

$$(\theta_1 - \theta_2)[t_1(x^u) - t_1(x^v)] + \sum_{i=1}^2 [\theta_i u_i(a^u) + (1 - \theta_i) v_i(a^v)] \quad (1)$$

Thus, if the agents could commit to coordinate in each state on an action profile that belongs to a different cell in X , then there would be no upper bound on the stakes of the bet that they would want to sign: if $\theta_1 > (<) \theta_2$, they would set $t_1(x^u) \gg (<) t_1(x^v)$. However, because the agents cannot commit, they must take into account their ability to manipulate the bet's outcome.

For instance, suppose that in the team example the agents agree on a bet that satisfies

$$t_1(x^l) - t_1(x^h) > u_1(h, h) - u_1(l, h)$$

Agent 1 would then choose l (exert no effort) in both states because the amount he gains in side payments outweighs the “bare” loss he incurs by deviating from the state u Nash equilibrium. But if agent 1 always shirks, the project would never be completed, and the agents would not be able to bet. Thus, in order to be sustainable, a bet must provide the agents with incentives to play in each state an action profile that belongs to a different cell in the partition.

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 maximizes (1) subject to the constraints

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

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

We refer to the value of the objective function (1), 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. The constraints of the optimization problem, given by the above pair of inequalities, constitute the second-period incentive constraints (SPIC). These ensure that indeed no agent has any incentive to unilaterally manipulate the bet in neither of the states. Formally, the SPIC constraints imply that a^u and a^v are pure-strategy Nash equilibria of the modified games $G(u, t)$ and $G(v, t)$. Our first result establishes that the SPIC constraints rule out the infinite-bets problem.¹

Proposition 1. For any finite G , the CIE surplus is finite and well-defined.

Note that this result relies on the assumption that the verifiability structure given by the partition X is state-independent. When the partition varies with the state, it is possible in some cases to sustain infinite bets.²

To better understand the notion of CIE surplus, let us apply it to our team example. Note first that because the bet assigns exactly the same transfers to (l, l) , (l, h) and (h, l) , the SPIC constraints imply that in each state ω , $a^\omega \in \{(l, l), (h, h)\}$.

Suppose $a^v = (h, h)$. Then the SPIC constraints in this state imply that

$$v_1(l, h) - v_1(h, h) \leq t_1(x^h) - t_1(x^l) \leq v_2(h, h) - v_2(h, l)$$

in contradiction to our assumption that l is a dominant strategy for each agent. It follows that $a^v = (l, l)$. Note that in this case, no agent has any incentive to deviate *no matter what bet is signed*: choosing h does not increase the bare game payoff and does not alter the transfers.

If $a^u = (l, l)$, then there will be no bets and the interim-expected surplus would be equal to $\sum_i [\theta_i u_i(l, l) + (1 - \theta_i) v_i(l, l)]$. However, we can attain a strictly higher surplus if we allow the agents to bet by letting $a^u = (h, h)$. The SPIC constraints in this state imply that

$$u_1(l, h) - u_1(h, h) \leq t_1(x^h) - t_1(x^l) \leq u_2(h, h) - u_2(h, l)$$

¹ If the modified game $G(\omega, t)$ does not have a pure-strategy NE, 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.

² We thank a referee for this point.

Since (h, h) is a Nash equilibrium in $G(u)$, $u_2(h, h) - u_2(h, l) \geq 0$ while $u_1(l, h) - u_1(h, h) \leq 0$. Thus, a tuple $((h, h), (l, l), t)$, where t satisfies the above inequality, attains a total interim-expected surplus, which equals

$$(\theta_1 - \theta_2)[u_2(h, h) - u_2(h, l)] + \sum_i [\theta_i u_i(l, l) + (1 - \theta_i)v_i(h, h)]$$

if $\theta_1 > \theta_2$, and it is equal to

$$(\theta_2 - \theta_1)[u_1(h, h) - u_1(l, h)] + \sum_i [\theta_i u_i(l, l) + (1 - \theta_i)v_i(h, h)]$$

if $\theta_2 > \theta_1$. Notice that each of these expressions is strictly higher than $\sum_i [\theta_i u_i(l, l) + (1 - \theta_i)v_i(l, l)]$. Therefore, $((h, h), (l, l), t)$ is CIE and the CIE surplus is given by the above expressions.

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 $\theta_i, \theta_j, \theta'_j$, $p_i(u | \theta_i, \theta_j) = p_i(u | \theta_i, \theta'_j)$. 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.³

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

2.1. 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 θ , there exists a bet $t(\theta)$ such that $[a^u, a^v, t(\theta)]$ is CIE for θ . In this case, we say that $[a^u, a^v, t(\theta)]$ is a purely speculative CIE tuple and that $t(\theta)$ is a *purely speculative CIE bet*.

³ However, Yildiz (2007) argues that 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.

Note that in general, attaining the CIE surplus may require a^u and a^v to vary with θ . 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 Nash equilibrium of the bare game. For instance, the CIE surplus in our team example is attained by pure speculation. To see why, recall that in the previous subsection we showed that the CIE surplus is attained by a bet that induces each agent to play his unique dominant strategy of the bare game in state v and to coordinate on the Pareto superior Nash equilibrium of the bare game in state u .

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 $d_i(a^\omega \rightarrow x)$ 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:

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

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

For the final piece of notation, let

$$D_1(a^u, a^v) \equiv \min_{y \in X} d_1(a^u \rightarrow y) + d_2(a^v \rightarrow y)$$

$$D_2(a^u, a^v) \equiv \min_{y \in X} d_2(a^u \rightarrow y) + d_1(a^v \rightarrow y)$$

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

When the underlying game 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 2. *A purely speculative CIE bet t satisfies the following property for every pair of priors θ :*

$$t_1(x^u) - t_1(x^v) = D^*(a^u, a^v | \theta) \quad (2)$$

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 $D^*(a^u, a^v | \theta)$, consider two special cases. First, in our team example, when both agents exert *no* effort in state v (i.e., $a^v = (l, l)$) the project is not completed (i.e., $x^v = x^l$), and neither agent can unilaterally change this outcome. However, when both agents exert effort in state u ($a^u = (h, h)$) the project is completed ($x^u = x^h$), but by shirking, each agent can unilaterally prevent the project from being completed (i.e., induce the outcome x^u). Therefore,

$$D_1(a^u, a^v) = u_1(a_1^u, a_2^u) - u_1(a_1^v, a_2^u) = u_1(h, h) - u_1(l, h)$$

and

$$D_2(a^u, a^v) = u_2(a_1^u, a_2^u) - u_2(a_1^u, a_2^v) = u_2(h, h) - u_2(h, l)$$

By Proposition 2, a purely speculative bet in this example satisfies the following for every pair of priors:

$$t_1(x^u) - t_1(x^v) = \begin{cases} u_2(h, h) - u_2(h, l) & \text{if } \theta_1 > \theta_2 \\ u_1(l, h) - u_1(h, h) & \text{if } \theta_1 < \theta_2 \end{cases} \quad (3)$$

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 Nash equilibrium 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 $d_1(a^u \rightarrow x^v)$. 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 $d_1(x^u \rightarrow y)$. 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 .

To conclude this sub-section, we need to emphasize that our focus on purely speculative CIE bets means that in certain classes of situations, our analysis is vacuous. For instance, when $G(u)$ and $G(v)$ have a unique Nash equilibrium which happens to be the same, pure speculation implies no betting and therefore, if betting can enhance the agents' interim surplus at all, it must distort the game's outcome at least in one state and therefore cannot be purely speculative. The general problem of characterizing the class of games for which the CIE surplus is attained by pure speculation is difficult and left for future research.

2.2. Implementation of purely speculative CIE bets

We now turn to the question of whether the CIE surplus can be implemented when the parties' priors are private information. Before presenting the implementation problem we consider, it is first instructive to note how privately known priors could act as a barrier to mutually beneficial speculative bets. Consider our team example, and assume that $u_1(h, h) - u_1(l, h) = u_2(h, h) - u_2(h, l)$. Suppose that in period 1, the two agents play the following naïve mechanism: each agent guesses whether the second-period outcome will be x^u or x^v ; when *exactly* one agent guesses correctly, he receives $\frac{1}{2}[u_1(h, h) - u_1(l, h)]$ from the other party; otherwise, no payments are made in period 2. Since $\theta_1 \neq \theta_2$ with probability one, the two agents can always earn speculative gains, if the agent with the higher prior on u guesses x^u , while the other guesses x^v . However, note that when $\theta_1, \theta_2 > \frac{1}{2}$, both agents would want to guess x^u . Similarly, when $\theta_1, \theta_2 < \frac{1}{2}$, both would want to guess x^v . 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.

Note that in the naïve mechanism whenever $\theta_i > \frac{1}{2}$, agent i would want to behave as if his prior is 1, and when $\theta_i < \frac{1}{2}$, he would want to behave as if it is 0. Thus, there is no predetermined direction in which an agent would want to exaggerate his private information. This contrast with a standard problem of trade between a buyer and a seller whose valuations of an asset are private information: the buyer would always want to exaggerate his value *downwards*, while the seller would want to exaggerate his value *upwards*. This suggests that the effect of asymmetric information in our model is similar to a situation in which two parties with private valuations of an asset wish to allocate this asset between them, but neither has full ownership rights over it (hence, each party is in some sense both a seller and a buyer). This analogy shall play an important role in our main result of this section.

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 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 $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2)$ is assigned a bet $t(\hat{\theta})$, and the agents play $G(\omega, t(\hat{\theta}))$ in period 2. Thus, a direct mechanism $t(\hat{\theta})$ induces a two-stage game with incomplete information, denoted $\Gamma(t)$.

Define

$$T_i^u(\theta'_i) \equiv \int_0^1 t_i(x^u \mid \theta'_i, \theta_j) dF(\theta_j)$$

and

$$T_i^v(\theta'_i) \equiv \int_0^1 t_i(x^v \mid \theta'_i, \theta_j) dF(\theta_j)$$

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 $t(\hat{\theta})$.

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

(EFF) $t(\hat{\theta})$ satisfies (2),

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 Nash equilibrium 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 $t(\hat{\theta})$ is a CIE bet. Condition PS-SPIC means that in the second stage of $\Gamma(t)$, the agents play a Nash equilibrium 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.⁴

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

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

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

$$q_i(\theta) = \begin{cases} r_1 + r_2 & \text{if } \theta_i \geq \theta_j \\ 0 & \text{if } \theta_i < \theta_j \end{cases}$$

(IC*) There is a Bayesian Nash equilibrium 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 Nash equilibrium 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 :

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

⁴ Note that the IR constraint captures the fact that the assumption of pure speculation implies that when the players do not sign a bet, they play the Nash equilibrium of the bare game.

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 3. *Let $\langle(u, v), G, X, F\rangle$ be a bilateral speculation problem for which the CIE surplus is attained by pure speculation for all profiles of priors, and which 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 Cramton et al. (1987), 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 for all profiles of priors. 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 these results, suppose that we can ignore the possibility that agents can manipulate the bet's outcome into some $y \neq x^u, x^v$. In this case:

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

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 Nash equilibria 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 Cramton et al. (1987) show to be implementable for any F . Thus, symmetric speculation problems occupy a special place in our model.

Let us use our team example to illustrate the intuition for Proposition 3 and Corollary 1. Suppose the agents use a mechanism that satisfies (3) but with regards to the agent's *reported* priors $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2)$. Then, regardless of the first-period outcome, there is a Nash equilibrium in the second-period subgame in which the agents play (h, h) in state u and (l, l) in state v . Moreover, if $\hat{\theta} = \theta$, then this mechanism assigns a CIE bet to every pair of priors. The problem is to design such a mechanism, which also ensures that the parties participate and report their true priors.

To tackle this problem, we reinterpret it as a standard allocation problem. Suppose that both agents report their true priors in period 1. Let $a^u = (h, h)$ and $a^v = (l, l)$ and denote $\delta_1(h, l) \equiv u_1(h, h) - u_1(l, h)$ and $\delta_2(h, l) \equiv u_2(h, h) - u_2(h, l)$. Consider the decision problem that an agent, say agent 1, faces in the second period. *What is agent 1's gain in period 2 from choosing a_1^ω in state ω relative to choosing l ?* By definition, the gain is zero in state v , regardless of whether $\hat{\theta}_1$ is higher or lower than $\hat{\theta}_2$. However, in state u the gain is $\delta_1(h, l) - [t_1(x^v | \hat{\theta}) - t_1(x^u | \hat{\theta})]$. By our construction of $t(\hat{\theta})$ and the assumption that $\hat{\theta} = \theta$, this difference is equal to zero when $\theta_2 > \theta_1$ and equal to $\delta_1(h, l) + \delta_2(h, l)$ when $\theta_1 > \theta_2$.

Thus, the agent's gain may be interpreted as a right to receive a prize of $\delta_1(h, l) + \delta_2(h, l)$, conditional on (h, h) being played in period 2. Put differently, the right is an asset of size $\delta_1(h, l) + \delta_2(h, l)$, whose first-period valuation by each agent is $\theta_i[\delta_1(h, l) + \delta_2(h, l)]$. Note that agent 1 receives this asset if and only if $\theta_1 > \theta_2$. This is analogous to

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.

allocating the asset to the person who values it the most. If no bet is signed in period 1, and the agents coordinate on (h, h) in state u , then agent 1's gain from choosing a_1^ω relative to choosing l is zero in state v and $\delta_1(h, l)$ in state u . Thus, it is as if agent 1 initially holds a share of $\delta_1(h, l)$ in the asset. His first-period valuation of this asset is $\theta_1 \delta_1(h, l)$. By signing the bet with the speculator, the agent increases his share by $\delta_2(h, l)$, as long as $\theta_1 > \theta_2$.

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 agents jointly hold an asset of size $\delta_1(h, l) + \delta_2(h, l)$. The agents' shares in the asset are $\delta_1(h, l)$ and $\delta_2(h, l)$. Each agent privately and independently draws a valuation of the asset. The problem is to design a mechanism that allocates the entire asset to the agent with the highest valuation, subject to the constraint that both agents agree to participate in this mechanism. Cramton et al. (1987) showed that implementing this objective depends on the initial ownership structure. When $\delta_1(h, l) \gg \delta_2(h, l)$ —that is, if agent 1 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 $\delta_1(h, l)$ and $\delta_2(h, l)$ shrinks, each agent 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 even in state u .

3. 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.

3.1. Bertrand competition

In this subsection, 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 4. *In the above bilateral speculation problem, the CIE surplus is attained by pure speculation for any profile of priors. Moreover:*

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

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

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

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

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.

3.2. Bilateral trade

In this subsection, 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 5. *In the above bilateral speculation problem, the CIE surplus is attained by pure speculation for any profile of priors. 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:*

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

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^h) are bare-game Nash equilibria in states l and h , respectively. Also note that part (ii) in Proposition 5 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 Nash equilibria. 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 .

4. 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 our team example, the first property holds whenever $u_1(h, h) - u_1(l, h) = u_2(h, h) - u_2(h, l)$, while the second property holds automatically because X has only two cells.

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 Cramton et al. (1987) to the language of our model.

In the Bertrand example of 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 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 team example, the mechanism means that the agents play a first-price auction in order to determine which of them wins the right to a prize of $\delta_1(h, l)$, conditional on the project being completed 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 whether the project is completed in the team example), 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 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 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:

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

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 θ), 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 1 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 our team example there were $n > 2$ team members. As in our original example, unless all members exert effort, the project is not completed. All other features of our example remain the same: the bare game in state u has a Pareto superior Nash equilibrium in which all members exert effort, in state u it is a dominant strategy for each agent to exert no effort, and agents can only bet on whether or not the project is completed. With slight abuse of notation, let $u_i(x^h)$ and $v_i(x^h)$ denote agent i 's bare-game payoff when the project is completed in state u and v respectively. Similarly, let $u_i(x^l)$ and $v_i(x^l)$ denote i 's payoff when the project is not completed in states u and v respectively.

The structure of CIE bets in this variant of the example is such that agent $i^* = \arg \min_i \theta_i$ —i.e., the agent who has the biggest faith in the completion of the project—essentially signs a bilateral side bet with every other agent. The stakes of the bilateral bet between i^* and $i \neq i^*$ are equal to $u_i(x^h) - u_i(x^l)$, namely i 's cost of interfering with the project's completion in state u by choosing l while all other agents choose h . 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 $\sum_i [u_i(x^h) - u_i(x^l)]$, in which the share of partner i in the jointly owned asset is $u_i(x^h) - u_i(x^l)$. Thus, using Propositions 1–3 in Cramton et al. (1987), it can be shown that as the utility differences $u_i(x^h) - u_i(x^l)$ 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 Spiegel (2006, 2008), 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 Spiegel (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 Spiegel (2008) 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 Cramton et al. (1987) was taken up by Fieseler et al. (2003) and Jehiel and Pauzner (2006), 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.

Acknowledgments

We thank the Editor, an Associate Editor and two anonymous referees whose comments helped improve this paper. We also wish to thank Barton Lipman, Eric Maskin, Ady Pauzner, Wolfgang Pesendorfer, Ariel Rubinstein and especially Eddie Dekel, for helpful conversations and comments. The first author acknowledges financial support from NSF Grant No. 0611938, and the second author is grateful for financial support from the US-Israel Binational Science Foundation, Grant No. 2002298.

Appendix A. Proofs

Proof of Proposition 1. Consider some CIE tuple (a^u, a^v, t) . By the SPIC constraints and budget-balancedness,

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

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 G is finite, $\theta_i u_i(a^u) + (1 - \theta_i) v_i(a^v)$ 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. \square

Proof of Proposition 2. Assume that the CIE surplus is attained by pure speculation and consider some purely speculative CIE tuple $(a^u, a^v, t'(\theta))$. Then for all $\theta = (\theta_1, \theta_2)$, the bet $t'(\theta)$ maximizes

$$(\theta_1 - \theta_2) \cdot [t'_1(x^u) - t'_1(x^v)] \quad (4)$$

subject to the SPIC constraints.

We proceed in two steps. First, we show that we can construct a bet t that satisfies (2) as well as the SPIC constraints. Second, we show that (2) 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 for all $y \in X$ either

$$t_1(y) - t_1(x^v) = \min[d_1(a^v \rightarrow y), D_2(a^u, a^v) + d_1(a^u \rightarrow y)] \quad (5)$$

or

$$t_1(y) - t_1(x^v) = \min[d_1(a^v \rightarrow y), -D_1(a^u, a^v) + d_1(a^u \rightarrow y)] \quad (6)$$

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 $y \in X$:

$$t_1(y) - t_1(x^u) \leq d_1(a^u \rightarrow y) \quad (7)$$

$$t_1(y) - t_1(x^v) \leq d_1(a^v \rightarrow y) \quad (8)$$

$$t_1(x^v) - t_1(y) \leq d_2(a^v \rightarrow y) \quad (9)$$

$$t_1(x^u) - t_1(y) \leq d_2(a^u \rightarrow y) \quad (10)$$

Suppose that $d_1(a^u \rightarrow y) + \hat{D} \leq d_1(a^v \rightarrow y)$. Then, by (5) and (6), inequalities (7) and (8) are satisfied. Assume that (9) is violated. Then, by (5) and (6):

$$-\hat{D} > d_1(a^u \rightarrow y) + d_2(a^v \rightarrow y) \quad (11)$$

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 (11) cannot exceed its R.H.S., a contradiction.

If $\hat{D} = D_2(a^u, a^v)$, then by our assumption that $[a^u, a^v, t'(\theta)]$ is a purely speculative CIE tuple,

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

contradicting (11). Therefore, (9) must hold. Finally, to see that (10) is satisfied, note that the L.H.S. of this inequality is equal to $-d_1(a^u \rightarrow y)$ and by our pure speculation assumption,

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

Alternatively, suppose that $d_1(a^u \rightarrow y) + \hat{D} > d_1(a^v \rightarrow y)$. Then, by (5), inequalities (7) and (8) are satisfied. Assume that (10) is violated. Then, by (5):

$$\hat{D} > d_1(a^v \rightarrow y) + d_2(a^u \rightarrow y) \quad (12)$$

If $\hat{D} = D_2(a^u, a^v)$, then by definition, it cannot exceed the R.H.S. of (12), a contradiction. If $\hat{D} = -D_1(a^u, a^v)$, then by our assumption that $[a^u, a^v, t'(\theta)]$ is a purely speculative CIE tuple,

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

contradicting (12). Therefore, (10) must hold. Finally, (9) follows from our pure speculation assumption, which implies that

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

This concludes the proof of the lemma. \square

Construct a bet t that satisfies (5) and (6) 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 $\hat{D} = D^*(a^u, a^v | \theta)$. Then for $y = x^u$, the bet t satisfies (2). 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:

$$-D_1(a^u, a^v) \leq t_1(x^u) - t_1(x^v) \leq D_2(a^u, a^v) \quad (13)$$

The SPIC constraints, summarized by (7)–(10), imply that for every $y \in X$:

$$-d_1(a^u \rightarrow y) - d_2(a^v \rightarrow y) \leq t_1(x^u) - t_1(x^v) \leq d_2(a^u \rightarrow y) + d_1(a^v \rightarrow y)$$

But this boils down to (13). Therefore, (2) is necessary for constrained interim-efficiency. \square

Proof of Proposition 3. 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 $\hat{\theta}$,

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

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

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

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

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

and

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

These two inequalities imply that $(q(\hat{\theta}), m(\hat{\theta}))$ 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 $(q(\hat{\theta}), m(\hat{\theta}))$ 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 $\theta \in [0, 1]^2$, 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 2, t satisfies (5) and (6), without loss of generality.

Let $t(x | \hat{\theta})$ 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 $\hat{\theta}$:

$$t_i(x^v | \hat{\theta}) = m_i(\hat{\theta}) \quad (14)$$

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

$$t_1(y | \hat{\theta}) - t_1(x^v | \hat{\theta}) = \min[d_1(a^v \rightarrow y), d_1(a^u \rightarrow y) + q_1(\hat{\theta}) - D_1(a^u, a^v)] \quad (15)$$

Because $(q(\hat{\theta}), m(\hat{\theta}))$ satisfies (EFF*), $q_1(\hat{\theta}) - D_1(a^u, a^v) = D^*(a^u, a^v | \theta)$. In particular, this means that $q_1(\hat{\theta}) - D_1(a^u, a^v)$ 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 $D_2(a^u, a^v)$, $d_1(a^u \rightarrow x^u) + q_1(\hat{\theta}) - D_1(a^u, a^v)$ cannot exceed $d_1(a^v \rightarrow x^u)$. Hence,

$$t_1(x^u | \hat{\theta}) - t_1(x^v | \hat{\theta}) = q_1(\hat{\theta}) - D_1(a^u, a^v) \quad (16)$$

The observation that $q_1(\hat{\theta}) - D_1(a^u, a^v) = D^*(a^u, a^v | \theta)$ implies that Eq. (16) becomes Eq. (2). This means that when $y = x^u$, $t(x | \hat{\theta})$ satisfies (EFF). By Lemma 1, this also means that $t(x | \hat{\theta})$ satisfies (PS-SPIC). It remains to show that $t(x | \hat{\theta})$ satisfies (IC) and (IR). Since $(q(\hat{\theta}), m(\hat{\theta}))$ 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 $t(x | \hat{\theta})$ and $q_1(\hat{\theta})$ given by (16), the last two inequalities imply (IC) and (IR), respectively. \square

Proof of Proposition 4. 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

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

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 inequality implies that $p^\omega = c^\omega$, a contradiction.

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

$$|\theta_1 - \theta_2| \cdot (c^H - c^L) \tag{17}$$

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,

$$c^L - c^H \leq t_1^L - t_1^H \leq c^H - c^L \tag{18}$$

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 $t(x | \theta)$ 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:

$$p_* - c^* - s_1 \cdot (p^* - c^*) \leq t_1(p^*) - t_1(p_*) \leq s_2 \cdot (p^* - c^*) + c^* - p_* \quad (19)$$

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 (19) imply that total surplus is less than (17). 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 (17), we complete the proof.

Suppose that $\theta_2^* \geq \theta_1^*$. Then, by (19), 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 (19), 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

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

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 (17). \square

Proof of Proposition 5. 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:

$$\begin{aligned} p^\omega - c + t_s(p^\omega) &\geq t_s(NT) \\ \omega - p^\omega - t_s(p^\omega) &\geq -t_s(NT) \end{aligned} \quad (20)$$

Hence,

$$c - p^\omega \leq t_s(p^\omega) - t_s(NT) \leq \omega - p^\omega \quad (21)$$

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

$$\pi_s(p^\omega - c) + \pi_b(\omega - p^\omega) + (\pi_s - \pi_b) \cdot [t_s(p^\omega) - t_s(NT)] \quad (22)$$

where π_i denotes agent i 's prior on ω .

By (21), 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:

$$\max\{\pi_b, \pi_s\} \cdot (h - c) \quad (23)$$

Now suppose that trade occurs in both states. Then the inequalities in (20) are the SPIC constraints that prevent each agent from unilaterally imposing no trade in each state. Hence,

$$p^l - p^h + c - l \leq t_s(p^h) - t_s(p^l) \leq p^l - p^h + h - c \quad (24)$$

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 (24), total surplus is bounded from above by:

$$\max\{\theta_b, \theta_s\}(h - c) + (1 - \min\{\theta_b, \theta_s\})(l - c) \quad (25)$$

which is below (23), since $l < c$. It follows that the value of the CIE surplus is given by (23), 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 (1), we obtain (23). 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. \square

References

- Allaz, B., Vila, J.-L., 1993. Cournot competition, forward markets and efficiency. *J. Econ. Theory* 59, 1–16.
- Bazerman, M.H., Gillespie, J.J., 1999. Betting on the future: The virtues of contingent contracts. *Harvard Bus. Rev.*, 155–160.
- Cramton, P., Gibbons, R., Klemperer, P., 1987. Dissolving a partnership efficiently. *Econometrica* 55, 615–632.
- Eliaz, K., Spiegler, R., 2006. Contracting with diversely naive agents. *Rev. Econ. Stud.* 73, 689–714.
- Eliaz, K., Spiegler, R., 2007. A mechanism-design approach to speculative trade. *Econometrica* 75, 875–884.
- Eliaz, K., Spiegler, R., 2008. Consumer optimism and price discrimination. Working paper, Brown University and Tel Aviv University.
- Fieseler, K., Kittsteiner, T., Moldovanu, B., 2003. Partnerships, lemons and efficient trade. *J. Econ. Theory* 113, 223–234.
- Jackson, M., Wilkie, S., 2005. Endogenous games and mechanisms: Side payments among players. *Rev. Econ. Stud.* 72, 543–566.
- Jehiel, P., Pauzner, A., 2006. Partnership dissolution with interdependent values. *RAND J. Econ.* 22, 1–22.
- Landier, A., Thesmar, D., in press. Financial contracting with optimistic entrepreneurs. *Rev. Finan. Stud.*
- Myerson, R., Satterthwaite, M., 1983. Efficient mechanisms for bilateral trading. *J. Econ. Theory* 29, 265–281.
- Neeman, Z., 1999. Property rights and efficiency of voluntary bargaining under asymmetric information. *Rev. Econ. Stud.* 66, 679–691.
- Newbery, D.M.G., 1984. Manipulation of futures markets by a dominant producer. In: Anderson, R.W. (Ed.), *The Industrial Organization of Futures Markets*. Lexington Books, Lexington, MA.
- Wilson, R., 1968. The theory of syndicates. *Econometrica* 36, 119–132.
- Yildiz, M., 2007. Wishful thinking in strategic environments. *Rev. Econ. Stud.* 74, 319–344.