Disagreement, Information, and Trade

Erik Eyster — UCSB

Tristan Gagnon-Bartsch — University of Iowa

Matthew Rabin* — Harvard University

September 4, 2026

Abstract

Asymmetries in information among rational traders are widely recognized as unable to explain observed levels of speculative trade in financial markets. Heterogeneous priors—in which agents "agree to disagree" about the likelihoods of different events, even contingent on the same information—have been posited to explain such trade. In a contractually rich, zero-sum setting, we explore some of the patterns relating trade generated by heterogeneous priors to the distribution of traders' information. Fixing an information structure, virtually any trade can be supported by some heterogeneous priors, provided no trader knows that they can only lose from the trade. However, the comparative-static patterns when fixing traders' prior beliefs and varying the assignment of information are far more restrictive. We establish a series of results showing that gainful trade flows in the direction of disagreements rather than private information. For instance, when traders' information is swapped, the signs of a mutually agreeable trade can reverse only if traders are not taking maximal advantage of their disagreement; in some cases, reversing the signs of mutually agreeable trade is impossible altogether. Thus, heterogeneous priors on their own cannot rationalize traders systematically trading in the direction of randomly assigned information.

*E-mails: erikeyster@ucsb.edu, tristan-gagnon-bartsch@uiowa.edu, and matthewrabin@fas.harvard.edu.

For helpful comments and suggestions, we thank Nageeb Ali and seminar participants at Berkeley, Bonn, Boston College, Cambridge, Capri "In Theory" conference, Harvard/MIT, Frankfurt School of Finance and Management and Goethe University, Harvard Business School, Iowa State, Penn State, Princeton, Queen Mary, St. Andrews, SAET Conference, Stanford, Toronto, UCL, UCLA Anderson, UCSB, UCSD, UNSW, UPenn, USC Conference on Social and Behavioral Finance, Washington University (Olin), and Yale SOM. We thank the Centre for Experimental Research on Fairness, Inequality and Rationality (FAIR) at the Norwegian School of Economics (NHH) for financial support and hospitality while conducting research on this paper.

1 Introduction

It is widely accepted that non-speculative motives alone cannot account for either the patterns or the volume of trade seen in real-world asset markets. The prediction of scant trade in asset markets reflects the logic of Tirole (1982) and Milgrom and Stokey (1982) that when agents sharing a common prior start from an efficient allocation, the arrival of additional information does not generate trade. Both because the assumption that differences in beliefs derive solely from differences in information seems unrealistic, and because the volume of trade is too high to be consistent with common-prior models, researchers have developed *non*-common-prior models: even people with the same information may disagree about the value of an asset. Some people may find the latest advice from a prominent financial guru compelling; others don't. When these heterogeneous priors are common knowledge—investors "agree to disagree"—would-be zero-sum trades become subjectively positive sum, and investors make trades in assets that allow them to bet against one another, all expecting to make profits.[1]

Although heterogeneous priors can give rise to speculative trade in a variety of settings, our central exercise is one of comparative statics: holding traders' prior beliefs—and hence the nature of their disagreement—fixed, how does changing the distribution of information affect the existence and direction of mutually agreeable trade? This exercise helps clarify what heterogeneous priors alone can explain when traders continue to draw rational inferences from others' willingness to trade.

We find that the direction of trade is substantially tied to the nature of disagreement rather than to the assignment of information. Concretely, trades that are mutually agreeable under one information assignment largely cannot be reversed under another. For example, although changing who knows what may alter the feasibility and magnitude of trade, it cannot transform a setting where Trader A bets against Trader B in favor of some event S into one in which Trader B bets against Trader A in favor of S. Heterogeneous priors on their own cannot rationalize retail investors systematically trading in the direction of randomly assigned information.

We analyze a stylized zero-sum environment with two risk-neutral traders. A trade specifies a transfer between them in each state of nature, and the traders have potentially different priors over those states.[2] Their information is represented by commonly known information partitions of the state space. Following Milgrom and Stokey (1982), Sebenius and Geanakoplos (1983), Geanakoplos (1994), and Morris (1994), we study which trades the traders would find strictly beneficial at every information set and call these trades *mutually agreeable*. This approach isolates the restrictions imposed by beliefs and information without requiring a complete specification of liquidity constraints, bargaining, prices, or market structure.

[1] Of course, some or all of them must be wrong on average.

[2] This is the "betting case" in the terminology of Morris (1994) and Gizatulina and Hellman (2019).

We begin our analysis with a basic result: for a given information structure, any trade that satisfies the minimal rationality assumption that no trader ever knows for sure that she is giving away money can be rationalized by some heterogeneous priors. In this static sense, heterogeneous priors can therefore rationalize almost anything. This observation motivates our analysis of the restrictions that emerge once we require the same priors to explain trade across different assignments of information.

Previous work identifies an important necessary condition for mutually agreeable trade. Wilson (1978) and Morris and Skiadas (2000) show that the information structure must contain a cycle: a sequence of states connected alternately through the two traders' information sets that eventually returns to its starting point.[3] A cycle allows gains in one state to compensate each trader for losses in another, given what that trader can distinguish. We represent information structures in a new way, namely as undirected bipartite graphs, with states as edges joining the traders' information sets. The signs of mutually agreeable trades can then be represented by orientations of these *information graphs*. This representation allows us to use existing results in graph theory to relate the direction of trade to the direction of disagreement along the cycles that support it.

In Section 3, we hold traders' priors fixed and ask how permuting the assignment of information sources affects the nature of mutually agreeable trades. We show that no mutually agreeable trade can be completely reversed when information is swapped. Indeed, this conclusion extends to any change in the information structure that preserves traders' public information.

We then show stronger results for information structures that contain a unique cycle. In such cases, not only can a mutually agreeable trade never be exactly reversed by swapping information, but every mutually agreeable trade under the swapped information structure must go in the same direction as the initial trade along the cycle. That is, the direction of trade on the cycle is pinned down by the nature of the disagreement, and is invariant to who has which information.

Even in information structures without a unique cycle, we identify a particular environment in which our results are strongest. Consider the case where all trades are *binary*, consisting of a fixed transfer from one party to the other in one set of states, and another fixed transfer from the other party in all other states. Fixing priors and public information, all mutually agreeable binary trades on a given event must go in the same direction, no matter the information structure.

These conclusions require qualification in more general environments. When information structures contain multiple cycles and trades can specify unrestricted transfers across states, mutually agreeable trades under two information assignments can sometimes have opposite signs in every state. Such reversals depend, however, on trades that contain unnecessarily large transfers. We formalize this observation using a weak notion of efficiency. A trade is *interim inefficient* if each

[3] Wilson (1978) establishes the importance of cycles for heterogeneous preferences more generally, whether the differences in marginal rates of substitution arise from heterogeneous priors or from other sources. Morris and Skiadas (2000) show that rationalizable trade does not exist in acyclic information structures.

trader, at each information set, would weakly prefer another trade that weakly shrinks the transfers made in every state, with one trader's preference at one information set strict. We show that two trades that are both mutually agreeable and interim efficient cannot have opposite signs in every state under the original and swapped information assignments. The no-reversal conclusion therefore extends beyond unique-cycle and binary environments once we exclude trades containing mutually undesirable excess transfers.

We also provide a positive counterpart to these no-reversal results. When every state belongs to a cycle, and traders disagree along at least one cycle, there exist mutually agreeable, interim-efficient trades before and after an information swap that have exactly the same signs: swapping information only alters the magnitudes of trade, not its direction. More generally, even when some states lie outside cycles and some signs must change with the assignment of information, the signs can be preserved on the minimal portion of trade that generates mutual gains by itself.

These results limit the extent to which heterogeneous priors can explain trade that follows the assignment of information. If information sources are assigned independently of traders' priors, we should expect little systematic correlation between the content of the information a trader happens to receive and the direction she trades.

This observation distinguishes heterogeneous priors from a conceptually distinct departure from the classic no-trade logic: traders may fail to account for others' private information when drawing inferences from their willingness to trade.[4] Eyster and Rabin (2005) formalize such behavior through the notion of *cursed equilibrium*, and Eyster, Rabin, and Vayanos (2019) extend the concept to canonical asset-market settings closer to the context studied here.[5] We do not provide a full treatment of cursedness, but in Appendix A we adapt the idea to our setting to highlight its contrast with heterogeneous priors. Whereas private information tends to impede mutually agreeable trade under heterogeneous priors, cursedness reverses this logic: whenever a trader has private information in a state, there is cursed trade that favors that trader. Moreover, whenever the traders have different information, the directions of cursed trade are independent of the nature of disagreement in their priors. Thus, cursedness can generate the prediction that the direction of trade follows the assignment of information.

[4]Indeed, we suspect that neglecting the role of information in trade may contribute to overstating heterogeneous priors as a source of trade. If we begin observing traders only after they have received their information, we may attribute differences in their behavior to intrinsic differences of opinion rather than to the happenstance of which information they received. If we peer in on small investors after they have completed their business-class flights, for example, we may interpret their different positions as reflecting different judgments rather than the different industry reports they happened to read along the way.

[5]There is evidence in laboratory settings of behavior arising from participants neglecting the information of others that cannot coherently come from heterogeneous priors (e.g., Carrillo and Palfrey, 2011), and the relationship between private information and directions of trade predicted by cursedness serves as an identifying assumption in empirical work such as Malmendier and Shanthikumar (2007). Fong, Lin, and Palfrey (2025) and Cohen and Li (2026) propose sequential-equilibrium-style refinements for cursed equilibrium in dynamic games.

Although our baseline setting of bilateral trade is clearly an abstraction, our underlying logic extends beyond that setting. In Section 4, we model two small investors embedded in a market with a third, "institutional trader" whom both expect to make money. Comparing the two small investors' trades essentially reduces to the bilateral case. We derive a similar comparison for a positive-sum environment such as demand for a medical treatment. Fixing patients' priors, swapping their information sources cannot switch which patient is commonly known to use a treatment and which patient is commonly known to refrain.

The paper's broader message is that heterogeneous priors do not provide an unrestricted account of speculative trade, once traders' priors are fixed across comparative statics. In the Conclusion, we discuss the extent to which our conclusions depend upon assumptions such as complete contracting and risk neutrality.

2 Cycles, Disagreement, and Trade

Two risk-neutral parties contemplate a zero-sum trade, which specifies how much money is transferred from Trader B to Trader A in each of a finite set of states of nature, E. We denote the trade by $t : E \to \mathbb{R}$. Trader A's utility from trade in $e \in E$ is $t(e)$, while Trader B's is $-t(e)$. We sometimes refer to subtrades of t that are non-zero only on a subset of states. For any $S \subseteq E$, let $t_S : E \to \mathbb{R}$ be defined by $t_S(e) = t(e)$ if $e \in S$ and $t_S(e) = 0$ otherwise.

Each trader has a prior belief over the set of states, E. We denote Trader i's prior distribution regarding $e \in E$ by $\mu^i \in \Delta(E)$, where $\mu^i(e) > 0$ for each $e \in E$ and each $i \in \{A, B\}$. Importantly, we do not assume that the traders share a common prior. For $S \subseteq E$, let $\mu^i(S) = \sum_{e \in S} \mu^i(e)$, and define conditional beliefs $\mu^i_S(e)$ such that $\mu^i_S(e) = \frac{\mu^i(e)}{\mu^i(S)}$ if $e \in S$ and $\mu^i_S(e) = 0$ otherwise. Additionally, define $w(e) := \ln \mu^A(e) - \ln \mu^B(e)$, which describes traders' difference of opinion regarding the likelihood of state e. Throughout the paper, we assume that traders' priors are fixed when performing comparative statics on information structures.

The two traders have partitional information about the state e. This means that each Trader i has a partition of E, $\mathscr{P}^i$, and learns which set element of $\mathscr{P}^i$ contains the true state e. We adopt a novel approach of representing the information structure as a labeled multigraph, which is a generalization of a labeled graph that allows for multiple edges between vertices.[6] This approach allows us to apply known results from graph theory to analyze speculative trade. Let $G = (V, E, \partial_0)$ be

[6] Rodrigues-Neto (2009, 2012) previously used graphs to represent information structures, and Hellwig (2013) offered a simplification of those results for some information structures. Our setup differs through different choices of edges and vertices. More recently, Cerreia-Vioglio, Corrao, and Lanzani (2025) employ a network structure similar to our graphical representation to study when a set of agents with non-subjective-expected-utility preferences have interim preferences that are consistent with a common ex-ante preference. This approach allows them to generalize Samet's (1998b) analysis on the existence of a common prior to the case of non-SEU preferences. See Golub and Morris (2017) for a further generalization of Samet (1998b).

a multigraph, where V is a finite set of vertices that represents information sets; E is a finite set of edges that represent states of nature; and $\partial_0 : E \to \{\{u, v\} \subset V : u \neq v\}$ is an incidence function that specifies the vertex endpoints of the edge e, thereby describing which information sets contain which states of nature. Thus, two vertices u and v being joined by edge e means that the information sets represented by those vertices both include state e.

The role assignment function $r : V \to \{A, B\}$ labels the vertices by assigning them to the two traders. The set $V^A := r^{-1}(A)$ collects A's information sets, while $V^B := r^{-1}(B)$ collects B's information sets. Recall that the function $w(e) = \ln \mu^A(e) - \ln \mu^B(e)$ assigns each edge a weight measuring belief disagreement. For simplicity, we abuse notation by writing $G = (V, E, \partial_0, r, w)$ and by referring to it as the *information graph*. Because each state belongs to exactly one information set for each trader, every edge joins one vertex in V^A to one in V^B. Thus, the information graph is bipartite: $\partial_0(e) = \{u, v\}$ implies that $r(u) \neq r(v)$. Let $E(v)$ be the set of edges incident to vertex v. When state $e \in E(v)$ occurs, the trader $r(v)$ learns that the state belongs to $E(v)$. This corresponds to each trader learning which set element of their information partition contains the state of nature. We assume that the information graph G is connected, which means that every pair of vertices is connected by a path. Equivalently, all states share the same smallest common-knowledge event, E.[7] Traders' priors and the information graph are common knowledge at every state.[8] At times, we refer to the subgraph of an information graph induced by the set of edges $S \subseteq E$. Let $\partial_0(S) = \cup_{e \in S} \partial_0(e)$ denote the set of vertices incident to the edges in S, and define the edge-induced information subgraph $G_S = (\partial_0(S), S, \partial_0|_S, r|_{\partial_0(S)}, w|_S)$.[9]

We frequently refer to the (real or notional) *ex ante* stage as the one before each trader receives their information, and the *interim stage* as the one after the traders learn their private information but before the state of nature is realized. Trader i uses Bayes' rule upon learning $e \in E(v)$ to form *interim beliefs*, $\mu^i_{E(v)} \in \Delta(E(v))$.

Some features of our setup are worthy of note. A main focus below revolves around comparative statics on the information graph, for instance positing that traders mutually agree to some trade t given the information partitions described by (V, E, ∂_0, r, w) and then asking which trades would be mutually agreeable given a role reversal in terms of information sources, $(V, E, \partial_0, r', w)$, where

[7]Formally, G is connected if for each pair of distinct vertices $v, v' \in V$, there exist edges $e_1, \ldots, e_{K-1} \in E$ and vertices $v_1, \ldots, v_K \in V$ such that $v_1 = v$, $v_K = v'$, and for each $k = 1, \ldots, K-1$, $\partial_0(e_k) = \{v_k, v_{k+1}\}$.

[8]For most of our results, these common-knowledge assumptions are less restrictive than they might first appear. Suppose that Trader B did not know whether Trader A possessed the vertex set V^A or $\widehat{V}^A$, and that this uncertainty were common knowledge. Then we could define a new state space to be $E \times \{1, 2\}$, where each new state specifies a state in the original edge set E and one of the two information structures for Trader A. Trader A's vertices would be $(V^A \times \{1\}) \sqcup (\widehat{V}^A \times \{2\})$, and Trader B's new information would remain V^B. Each state $(e, 1)$ would be incident to one vertex in $V^A \times \{1\}$ and one in V^B; each state $(e, 2)$ would be incident to one vertex in $\hat{V}^A \times \{2\}$ and one in V^B. By duplicating the set of edges, this approach enlarges the set of feasible trades. We further discuss the issue of contractual richness in the conclusion.

[9]The edge weights $w(e)$ are inherited from G and generally will not equal $\ln \mu^A_S(e) - \ln \mu^B_S(e)$. This does not affect any of our results.

$r'(v) \neq r(v)$ for each $v \in V$. That is, Trader A now has the partition B had in the initial case, and vice versa. Framed this way, it is clear that the states over which trades are defined do not include a description of the information graph itself; an alternative, and more complete, formulation of the notion of states sometimes employed would embed all the agents' information into the definition of the states themselves.[10] Additionally, unlike employment of similar frameworks used in the literature that refer solely to players' beliefs at the "interim stage" once players get signals determining their information sets, we include as a primitive players' prior beliefs over states "before" receiving information. Doing so lets us make sense of some of the exercises we perform below, most notably those regarding the interchange of players' information partitions.[11] And, as noted in the introduction and further below, our main analysis does not impose any restrictions on how t depends on states, allowing trade to depend on information that may be private at time of trade, and even more granular details of the environment.

Definition 1. *A set of edges $C \subseteq E$ is a* cycle *in $G = (V, E, \partial_0, r, w)$ if the edge-induced subgraph G_C is connected, and every vertex of G_C is incident to exactly two edges.*[12]

A cycle in our information structure corresponds to a walk along distinct edges and through distinct vertices that ends at its starting point. The walk must alternate between the two traders' information sets since the graph is bipartite. As a consequence, every cycle has an even number of edges. We refer to an information graph as *acyclic* when it contains no cycle and *cyclic* when it contains at least one cycle.

Before turning to the role of cycles in explaining speculative trade fueled by heterogeneous priors, as explored in the previous literature, it is useful to delineate a relationship between cycles and a minimal necessary criterion for trade. This relationship offers a more straightforward approach to proving some pre-existing and new results.

Even without specifying beliefs of the traders or determining whether a trade will be acceptable

[10]Beginning with Milgrom and Stokey (1982), many papers in the literature study how the refinement of private information affects trade, thereby implicitly assuming, like us, that the state does not encode traders' information. While all of our information-permuting results could be reformulated using an alternative framework that embeds the information partitions themselves into state space, doing so would lead to more cumbersome notation and characterization of the results. Under this alternative definition of states, we would consider an expanded state space where each state under our original definition is duplicated: one duplicate corresponds to players having one assignment of information partitions while the other corresponds to the reversed assignment. Players then have priors defined over this expanded state space. The equivalent of our information-permuting exercise is thus the following: if traders agree to some trade t that is non-zero only on states in which players have one assignment of information partitions, then what trades $\hat{t}$ will be accepted when $\hat{t}$ is non-zero only on states where players have the reversed assignment.

[11]If an information-belief structure only defined traders' beliefs on their own information sets, then we would not know what a trader would believe on other information sets, including those that belong to the other trader, which would prevent us from asking questions about behavior after transposing information partitions.

[12]Alternatively, a set of distinct edges $\{e_1, e_2, \ldots, e_K\} \subseteq E$ with $K \geq 2$ is a *cycle* in $G = (V, E, \partial_0, r, w)$ if K is even and there exist distinct vertices $\{v_1, v_2, \ldots, v_K\}$ such that $\partial_0(e_k) = \{v_k, v_{k+1}\}$ for $k = 1, \ldots, K$, where $v_{K+1} := v_1$.

to both parties, a necessary condition for each trader to strictly prefer the trade to no trade at each of their information sets is that the trader must benefit from the trade in at least one state in each of their information sets. The existence of trades meeting this test provides an alternative characterization of cycles. (Proofs of all results are presented in Appendix C).

Lemma 1. *The information graph $G = (V, E, \partial_0, r, w)$ contains a cycle if and only if there exists a trade t such that*

1. *for every A-vertex $v \in V^A$, there exists an incident edge $e \in E(v)$ with $t(e) > 0$; and*
2. *for every B-vertex $v \in V^B$, there exists an incident edge $e \in E(v)$ with $t(e) < 0$.*

As the proof of Lemma 1 makes clear, the result extends to non-zero-sum trades with opposing interests, where in every state at most one player benefits. This encompasses positive-sum-trade settings such as insurance markets.[13]

Turning to which trades rational traders betting on disagreements may engage in, we call a trade t *mutually agreeable at* e if it is common knowledge at e that both players strictly prefer executing t to no trade. Focusing on mutually agreeable trades allows us to sharply state our main points with simple arguments. Our assumption that the information graph is connected implies that the trade t is mutually agreeable at some e if and only if it is mutually agreeable at all e. Thus, in this case, we simply call t *mutually agreeable*. Restricting attention to connected information graphs comes with little loss of generality, since statements about mutual agreeability apply to each connected component across the graph.

Lemma 1 implies that mutually agreeable trade occurs exclusively in cyclic information structures. Previous work has explored how cycles in an information structure relate to the existence of a common prior, and thus to the possibility of speculative trade. Each trader has a system of interim beliefs $\{\mu^i_{E(v^i)}\}_{v^i \in V^i}$, one corresponding to each information set. Taking these interim belief systems as primitive, one can ask whether they could derive from a common prior, namely whether there exists some prior $\mu \in \Delta(E)$ with the property that for each $i \in \{A, B\}, v^i \in V^i$, and $e \in E(v^i)$,

$$\mu(e|E(v^i)) = \mu^i_{E(v^i)}(e).$$

Samet (1998a) shows that the two traders' interim belief systems are consistent with a common prior if and only if the intersection of their convex hulls is nonempty. Morris (1994) and Feinberg (1995) also provide necessary and sufficient conditions for the existence of a common prior. In other words, the traders' interim belief systems are consistent with a common prior if and only if

[13] Morris and Skiadas (2000) explore speculative trade, including non-zero-sum opposing-interest trade, for two players with heterogeneous priors, using the concept of interim rationalizability. Despite using a different solution concept and having a different focus, their Proposition 1 implies that an acyclic information structure does not generate trade. In general, rationalizable trade allows for more trading outcomes than the concept we explore below.

there exists some prior distribution that is a convex combination of each trader's interim beliefs. The existence of a common prior is equivalent to the impossibility of a mutually agreeable trade (Morris, 1994; Feinberg, 1995; Samet, 1998a).

Following this approach, Rodrigues-Neto (2009) takes interim beliefs as primitive and explores the role that cycles play for the existence of a common prior.[14] We can recast these previous results on the existence of a common prior in terms of cycles in the information graph.

Proposition 0. 1. *(Rodrigues-Neto, 2009, 2012; Hellman and Samet, 2012) If the information graph G is acyclic, then any pair of traders' priors (μ^A, μ^B) produces interim beliefs that are consistent with a common prior.*

2. *(Nyarko, 2010; Hellman and Samet, 2012) If the information graph G is cyclic, then the set of traders' priors (μ^A, μ^B) that produce interim beliefs consistent with a common prior has Lebesgue measure zero.*

Proposition 0 has immediate implications for the existence of speculative trade:

Corollary 0. 1. *If G is acyclic, then there exists no mutually agreeable trade.*

2. *If G is cyclic, then the set of traders' priors (μ^A, μ^B) that preclude mutually agreeable trade has Lebesgue measure zero.*

Proposition 0 summarizes results that have been established previously in the literature regarding the existence of a common prior. Because the existence of a common prior is equivalent to the absence of mutually agreeable trade, these results also offer necessary and sufficient conditions for speculation as presented in Corollary 0. Alternatively, a direct intuition for Part 1 of Corollary 0 comes from Lemma 1. The appendix contains a proof of Part 2 of Corollary 0 based on the approach developed in the remainder of this section of the paper.

To illustrate, consider a simple example in which $E = \{e_1, e_2\}$, and there is one A-vertex incident to both edges (see Figure 1). Consider two cases. In the first (Panel a), the two edges are incident to two distinct B-vertices. The graph is a tree, and hence acyclic: no mutually agreeable trade exists. In the second case (Panel b), the two edges are incident to a single B-vertex, and the graph consists of two parallel edges between an A-vertex and a B-vertex. This forms a cycle of length 2, and mutually agreeable trade is possible whenever priors differ.[15]

Before further exploring the nature of trade on cycles, recall that the introduction highlighted two primary motivations for moving beyond models with common priors. The first is to dispense

[14] See also Rodrigues-Neto (2012). Hellman and Samet (2012) provide a related approach—their notion of "tightness" is equivalent to acyclicity in our information graph—and discuss a weaker result in Harsanyi (1968).

[15] To see the connection between our graphical approach and the underlying information structures, Appendix B redraws all of the information graphs appearing in figures throughout the paper in terms of the underlying information partitions. See Figure B.1 for the partitional analog of Figure 1.

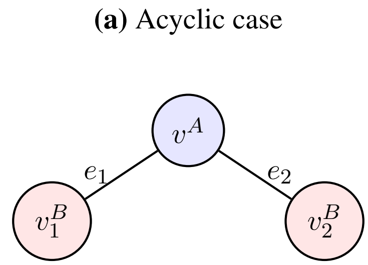

Figure 1. Two information structures with $E = \{e_1, e_2\}$ and a single A-vertex. Panel (a): the graph is acyclic, which precludes mutually agreeable trade. Panel (b): there is a cycle of length 2; mutually agreeable trade is possible whenever priors disagree.

with an unrealistic assumption in economic modeling. The second is to develop models that better reflect empirical evidence, especially the prevalence of speculative trade. Although our framework does not directly predict the level of trade, it does clarify which patterns of trade can be explained by heterogeneous priors. Most of our paper fixes traders' prior beliefs and explores how mutually agreeable trade depends upon traders' information. We begin, however, by characterizing mutually agreeable trades without any restrictions on priors.

Proposition 1 demonstrates that any pattern can be accounted for by heterogeneous priors, as long as it satisfies a minimal rationality criterion: no trader agrees to give away money with certainty, given their information.

Proposition 1. *There exist prior beliefs $\mu^A, \mu^B \in \Delta(E)$ for which t is mutually agreeable on the associated information graph $G = (V, E, \partial_0, r, w)$ with $w = \ln \mu^A - \ln \mu^B$ if and only if the following conditions hold:*

1. *for each $v^A \in V^A$, there exists an incident edge $e \in E(v^A)$ with $t(e) > 0$; and*
2. *for each $v^B \in V^B$, there exists an incident edge $e \in E(v^B)$ with $t(e) < 0$.*

Hence, even if we know traders' information, the only conclusion we can draw without knowing the specifics of their disagreements is that no party will willingly give away money.

Let us return to our main focus in which we fix the traders' prior beliefs. Above, we established that cycles in the information graph allow for mutually agreeable trade. We now further analyze features of agreeable trades, including which direction the transfers will flow. To begin, return to the two-state example in Figure 1. Suppose that

$$\frac{\mu^A(e_1)}{\mu^A(e_2)} > \frac{\mu^B(e_1)}{\mu^B(e_2)}, \tag{1}$$

or, equivalently, that

$$[\ln(\mu^A(e_1)) - \ln(\mu^B(e_1))] - [\ln(\mu^A(e_2)) - \ln(\mu^B(e_2))] = w(e_1) - w(e_2) > 0, \tag{2}$$

which, since there are only two states, means that Trader A assigns higher probability to e_1 than does Trader B. Any mutually agreeable trade t must have the property that $t(e_1) > 0 > t(e_2)$. The trade must follow the direction of disagreement: since Trader A assigns higher probability to e_1, Trader A must gain in e_1.

We now extend this logic to study the direction of trade on longer cycles. To do so, it is useful to describe both the direction of traders' disagreement and the direction of trade; we describe these directions relative to a fixed orientation of the information graph. We give G the simple reference orientation that orients every edge from its B-endpoint to its A-endpoint. We write the resulting directed labeled graph as $\vec{G} = (V, E, \partial_1, r, w)$, where, for an edge e from $v^B \in V^B$ to $v^A \in V^A, \partial_1(e) = v^A - v^B$.

Cycles are a necessary condition for trade. To characterize the direction of trade, we need to describe the pattern of trade along different cycles. To do so, we continue to work with the underlying undirected information graph G and record whether each edge is traversed along or against the reference orientation in $\vec{G}$.

Definition 2. *Let C be a cycle. An* orientation *of C is a function*

$$\sigma_C : C \to \{\pm 1\}$$

such that

$$\sum_{e \in C} \sigma_C(e)\, \partial_1(e) = 0.$$

An oriented cycle *$\vec{C}$ is a cycle C together with one of its two orientations.*

The orientation of a cycle pins down its direction of traversal. For example, in the right panel of Figure 1, the cycle $\{e_1, e_2\}$ can be oriented in two ways. In one orientation, $\sigma_C(e_1) = 1, \sigma_C(e_2) = -1$, and

$$\sum_{e \in C} \sigma_C(e)\, \partial_1(e) = \partial_1(e_1) - \partial_1(e_2) = v^A - v^B - (v^A - v^B) = 0;$$

in the other orientation, $\sigma_C(e_1) = -1, \sigma_C(e_2) = 1$, and

$$\sum_{e \in C} \sigma_C(e)\, \partial_1(e) = -\partial_1(e_1) + \partial_1(e_2) = -(v^A - v^B) + (v^A - v^B) = 0.$$

Under the first orientation, the oriented cycle goes from v^B to v^A and returns to v^B (counterclockwise); under the second orientation, it goes in the opposite direction (clockwise).

For an oriented cycle $\vec{C}$, define

$$W(\vec{C}) := \sum_{e \in C} \sigma_{\vec{C}}(e)\, w(e),$$

where C contains the edges of $\vec{C}$. The quantity $W(\vec{C})$ measures the direction and magnitude of the traders' disagreement around the oriented cycle. Reversing the orientation of $\vec{C}$ reverses the sign of $W(\vec{C})$. Rodrigues-Neto (2009, 2012) calls the set

$$\{W(\vec{C}) = 0 : \vec{C} \text{ is an oriented cycle in G}\} \tag{3}$$

the *cycle equations* and shows that they all hold if and only if the information structure is consistent with a common prior.

Definition 3. *Let $\vec{C}$ be an oriented cycle on states C. We say that a trade t* follows the direction of disagreement along the cycle $\vec{C}$ *if*

1. $W(\vec{C}) > 0 \Rightarrow \sigma_{\vec{C}}(e)t(e) > 0$ *for each* $e \in C$*;*
2. $W(\vec{C}) < 0 \Rightarrow \sigma_{\vec{C}}(e)t(e) < 0$ *for each* $e \in C$*.*

$W(\vec{C}) > 0$ means that the geometric average of Trader A's likelihood ratios of edges in the oriented cycle traversed towards A vertices to those traversed towards B vertices exceeds the corresponding geometric average for Trader B. A trade follows the direction of disagreement along a cycle when the signs of the trade follow the direction of traders' disagreement along the cycle. A trade *goes against the direction of disagreement* if it has the opposite signs from a trade that follows the direction of disagreement.

Whenever traders disagree along some cycle, there exists a mutually agreeable trade that follows the direction of disagreement along that cycle.

Lemma 2. *If for some oriented cycle $\vec{C}$, $W(\vec{C}) \neq 0$, then there exists a mutually agreeable trade that follows the direction of disagreement along $\vec{C}$.*

Disagreement along cycles creates the gains from trade that enable mutually agreeable trade.

Moreover, to take advantage of those gains from trade, a mutually agreeable trade must follow the direction of disagreement somewhere on the information graph. Proposition 2 makes this more precise.

Proposition 2. *Any mutually agreeable trade follows the direction of disagreement along some oriented cycle $\vec{C}$ for which $W(\vec{C}) \neq 0$.*

When an information structure has multiple cycles, it is possible that a mutually agreeable trade does not follow the direction of disagreement along *all* oriented cycles, and we will see below that it may even go against the direction of disagreement on some oriented cycle.[16]

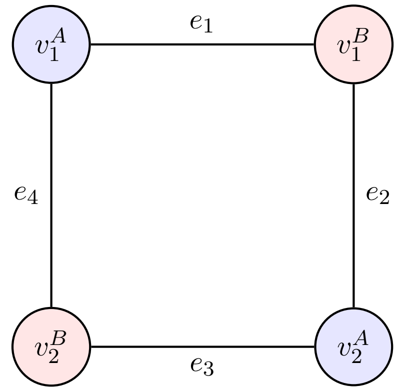

C_4 cycle

Figure 2. A four-state information structure forming a cycle.

Example 1. Figure 2 illustrates a single-cycle information structure. Consider $\mu^A = \left(\frac{3}{9}, \frac{1}{9}, \frac{2}{9}, \frac{3}{9}\right)$ and $\mu^B = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$ over $E = \{e_1, e_2, e_3, e_4\}$. Let $\vec{C}$ be the oriented cycle that goes clockwise:

$$\sigma_{\vec{C}}(e_1) = -1, \quad \sigma_{\vec{C}}(e_2) = +1, \quad \sigma_{\vec{C}}(e_3) = -1, \quad \sigma_{\vec{C}}(e_4) = +1;$$

$$\begin{aligned} W(\vec{C}) = \sum_{e \in C} \sigma_{\vec{C}}(e)\, w(e) &= -w(e_1) + w(e_2) - w(e_3) + w(e_4) \\ &= -\ln\left(\frac{3/9}{1/4}\right) + \ln\left(\frac{1/9}{1/4}\right) - \ln\left(\frac{2/9}{1/4}\right) + \ln\left(\frac{3/9}{1/4}\right) = -\ln 2. \end{aligned}$$

Trade follows the direction of disagreement only if

$$t(e_1), t(e_3) > 0 \qquad \text{and} \qquad t(e_2), t(e_4) < 0.$$

The trade $t = (6, -7, 4, -5)$, for instance, is mutually agreeable. Additionally, since there is a single cycle in the information graph, Proposition 2 further says that *any* mutually agreeable trade in this information structure has to have this sign pattern: A must be paid on the odd states whereas

[16] Indeed, it may be impossible to follow the direction of disagreement along all cycles. For example, when $E = \{e_1, e_2, e_3\}$, and A and B each possess one vertex, then any trade over the three states must have the same sign on at least two states, for example on e_1, e_2. Whenever $w(e_1) \neq w(e_2)$, the trade does not follow the direction of disagreement along the cycle on $\{e_1, e_2\}$. Neither does it go against the direction of disagreement on that cycle.

B must be paid on the even states. This foreshadows our results below showing that the direction of mutually agreeable trade has a weak relationship with the assignment of information sources. In this example, if we were to reverse the role assignment in the graph—which corresponds to swapping the information partitions of the two traders—then the terms in the calculations above would all flip signs and we would reach the same conclusion: $t(e_1), t(e_3) > 0$ and $t(e_2), t(e_4) < 0$. Thus, the direction of mutually agreeable trade in this example does not depend on the assignment of information sources.

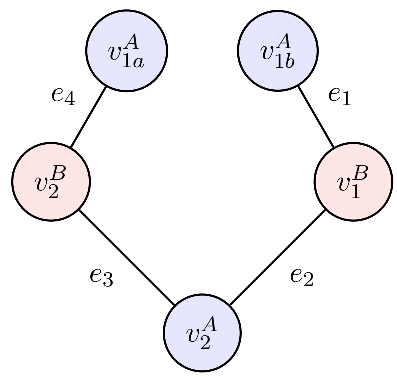

Figure 3. An acyclic information structure obtained by splitting v_1^A in Figure 2.

Disagreement creates the latitude for trade. What role does asymmetric information play? Let us return to Example 1, but suppose that Trader A can now distinguish state e_1 from e_4. This change splits Trader A's vertex v_1^A into two vertices, v_{1a}^A and v_{1b}^A, one of which is incident to e_1 and the other to e_4. Figure 3 illustrates this change. The new information structure is acyclic, which by Corollary 0 implies that there can be no mutually agreeable trade, regardless of priors.[17] Our next result clarifies how asymmetric information impedes trade in general. Before stating our result, we formalize the notion of refining traders' information in graph-theoretic terms.

Definition 4. *$\widehat{G} = (\widehat{V}, E, \widehat{\partial}_0, \widehat{r}, w)$ is* obtainable from $G = (V, E, \partial_0, r, w)$ by refining information *if there exists some surjection $\phi : \widehat{V} \to V$ such that for each $e \in E$, $\widehat{\partial}_0(e) = \{u, v\}$ implies that $\partial_0(e) = \{\phi(u), \phi(v)\}$, and for each $\widehat{v} \in \widehat{V}$, $\widehat{r}(\widehat{v}) = r(\phi(\widehat{v}))$.*

The mapping ϕ associates every vertex in the new graph with one in the original graph. Because the mapping is onto, every "parent" vertex in the original graph is preserved or gets split into multiple vertices. If a new vertex is incident to some edge, then its parent vertex is also incident to that edge. Moreover, the traders' labels and disagreement are preserved: refining information does not swap the traders' information sets or change their priors.

[17]Accordingly, the traders' interim beliefs are consistent with the common prior $\mu = \left(\frac{1}{6}, \frac{1}{6}, \frac{1}{3}, \frac{1}{3}\right)$. The results in Samet (1998b) imply that the common prior consistent with both traders' interim beliefs is unique.

In this proposition, we allow for the possibility that refining players' information transforms a connected graph into a disconnected graph.[18]

Proposition 3. *Let* $\widehat{G} = (\widehat{V}, E, \widehat{\partial}_0, \widehat{r}, w)$ *be obtainable from* $G = (V, E, \partial_0, r, w)$ *by refining information.*

1. *(Dow, Madrigal, and Werlang (1990), Gizatulina and Hellman (2019)) If there exists no* t *that is mutually agreeable under* G*, then there exists no* $\hat{t}$ *that is mutually agreeable under* $\widehat{G}$*.*

2. *If* t *is mutually agreeable under each connected component of* $\widehat{G}$*, then* t *is mutually agreeable under* G*.*

Part 1 of Proposition 3 states that refining players' information cannot create trade when none was previously possible, or, equivalently, that coarsening players' information cannot destroy the scope for trade. Dow et al. (1990) and Gizatulina and Hellman (2019) have previously stated and proved Part 1 for the special case, as expressed within our framework, where V has only two elements.[19] A simple intuition for Part 1 of Proposition 3 is that refining either player's information neither introduces new cycles nor affects disagreement along existing cycles. A simple intuition for Part 2 of Proposition 3 is that for any disjoint S_1, S_2,

$$\operatorname{sgn}(\mathbb{E}[t|S_1]) = \operatorname{sgn}(\mathbb{E}[t|S_2]) \Rightarrow \operatorname{sgn}(\mathbb{E}[t|S_1]) = \operatorname{sgn}(\mathbb{E}[t|S_2]) = \operatorname{sgn}(\mathbb{E}[t|S_1 \cup S_2])).$$

Put differently, if a trader's expected utility is positive at each information set with finer information, then it must also be positive at each information set with coarser information.

Our next result underscores that heterogeneous-priors-based trade does not depend upon asymmetric information: any trade that is mutually agreeable under asymmetric information is also mutually agreeable under some information structure with symmetric information. Thus, while the motivation to trade may be robust enough to overcome asymmetric information, such trades are driven purely by disagreement rather than by informational asymmetries.

Corollary 1. *If* t *is mutually agreeable under some* $\widehat{G}$ *obtainable from* $G = (V, E, \partial_0, r, w)$ *by refining information, where* V *has only two elements, then* t *is mutually agreeable under* G*.*

The condition that V has only two elements means that neither trader possesses any information. Hence, the scope for trade does not hinge upon either trader having private information.

[18] In the language of information partitions, we allow for the possibility that refining partitions leads to a non-singleton meet of the refined partitions.

[19] With risk-neutral traders, the absence of mutually agreeable ex ante trade implies that $\mu^A = \mu^B$, whereas the absence of mutually agreeable interim trade does not, which differentiates our result from those of Dow et al. (1990) and Gizatulina and Hellman (2019). The result of Dow et al. (1990) extends beyond expected-utility preferences. The result of Gizatulina and Hellman (2019) encompasses trade to which both parties agree in certain states, but not necessarily in all states.

3 Gainful Trade Does Not Flow Towards Private Information

In this section, we formalize some of the relationships between the structure of information and the flow of mutually agreeable trades. Imagine that Trader A reads news about a stock, and subsequently the two parties agree to a trade t where Trader A buys that stock from Trader B. The simplest result is that the reverse of a mutually agreeable trade $(-t)$ is not mutually agreeable when information is swapped. But we establish several results that strengthen this along two dimensions: first, we characterize the mutual agreeability of trades when the traders swap information—providing a more general description of how mutually agreeable trade can change—and second, we show how the mutual agreeability of a given trade holds up over a more general class of informational transformations.

We begin with an analysis of simple *binary* trades where we obtain a particularly strong result that also provides intuitions for our more general results. We define a trade t to be binary if its image $t(E) = \{x, y\}$ for some $x, y \in \mathbb{R}$ with $x \neq y$. Let $S = t^{-1}(x)$. The trade t is equivalent to a bet on whether event S occurs. Denote the space of all such binary bets on S by $T_{\text{BIN}}(S)$.[20]

Proposition 4. *1. Let t, t' be mutually agreeable binary trades under connected information graphs G, G' with common edge set E and edge weights w. Then it cannot be the case that for each $e \in E$, $\text{sgn}(t(e)) = -\text{sgn}(t'(e))$.*

2. Suppose that for some $S \subseteq E$, the trades $t, t' \in T_{\text{BIN}}(S)$ are mutually agreeable under connected information graphs G, G', respectively, which have common edge set E and edge weights w. Then for each $e \in E$, $\text{sgn}(t(e)) = \text{sgn}(t'(e))$.

The condition that the two information graphs G, G' are connected means that the two graphs have the same public information, or identical smallest common knowledge events. Traders' private information may vary between G and G', but the public information does not. One important example is when G' is obtained from G by swapping the traders' information partitions.

Two mutually agreeable binary trades cannot have opposite signs, regardless of the traders' information. Part 1 of Proposition 4 emphasizes that this holds even when the two trades are over different events. Furthermore, if t is a mutually agreeable bet on S under one information structure, then any mutually agreeable trade t' on S under any alternative information structure must have identical signs as t.

Turning from binary to more general trades, our results are not as simply stated nor quite as sharp. We first show that two parties who deem t to be mutually agreeable under the information graph will never agree to the reversed trade, $-t$, in any information structure that preserves public information. This includes the case where traders' information is swapped. We in fact show

[20] More specifically, $T_{\text{BIN}}(S) = \{t \in \mathbb{R}^{|E|} : \forall\, e, e' \in S,\ t(e) = t(e') \text{ and } \forall\, e, e' \in E \setminus S,\ t(e) = t(e')\}$.

a slightly stronger version of this: the traders also will not agree to the reversed trade plus a constant, which can be interpreted as changing the price at which traders exchange an asset with state-contingent payoffs.

Proposition 5. *Let t be mutually agreeable under the connected information graph G with edge set E, and let $\mathbf{1} \in \mathbb{R}^{|E|}$ denote a vector of ones. For any scalar $k \in \mathbb{R}$, the trade $-t + k\mathbf{1}$ is not mutually agreeable under any connected information graph G' that has the edge set E and edge weights w.*

Proposition 5 implies that two traders who agree on t under an original information structure cannot agree on the reversed trade, $-t$, when their information is permuted. It cannot happen, for instance, that whoever reads *Disagreement Today* trades shares in the direction suggested by its news articles, even if the stock prices could somehow depend upon who reads DT.

Proposition 6 describes more stringent limits on the scope for trade for the special case in which the information graph contains a unique cycle, as in Example 1. In this case, swapping the traders' information cannot flip all of the signs of the trade, and must preserve the signs on the cycle.

Starting from the information graph $G = (V, E, \partial_0, r, w)$, define the *swapped information graph* $G^\tau = (V, E, \partial_0, r^\tau, w)$, where r^τ swaps the traders' labels from those described by r, i.e., for each $v \in V, r(v) \neq r^\tau(v)$.[21]

To state our next result, we wish to distinguish between subtrades which are mutually beneficial and those which are not.

Definition 5. *Let $S \subseteq E$ and let t_S be a trade supported on S. We say that t_S is S*-mutually agreeable *if t_S is mutually agreeable on the edge-induced subgraph G_S.*

The concept of S-mutual agreeability captures the idea that the trade on the states in S generates mutual gains by itself, without relying on trade in states outside of S.

Proposition 6. *Let t be mutually agreeable under an information graph G with a single cycle C, and let t^τ be mutually agreeable under the swapped information graph G^τ. Then, for each $e \in C$,* $\text{sgn}(t(e)) = \text{sgn}(t^\tau(e))$*, while for each $e \in E \setminus C$,* $\text{sgn}(t(e)) = -\text{sgn}(t^\tau(e))$*. Moreover, t_C is C-mutually agreeable under G, and t^τ_C is C-mutually agreeable under G^τ. If $E \setminus C \neq \emptyset$, then neither $t_{E\setminus C}$ nor $t^\tau_{E\setminus C}$ is $E \setminus C$-mutually agreeable under G or G^τ.*

In the single-cycle case, not only can swapping the traders' information not reverse all of the signs of trade, but it must preserve the signs of trade along the cycle. Disagreement fully determines the direction of trade along the cycle, and the assignment of information plays no role. The assignment of information does affect the signs of trade off the cycle. However, because the

[21]Note that G and G^τ are obtainable by refining information from a common two-vertex graph on E.

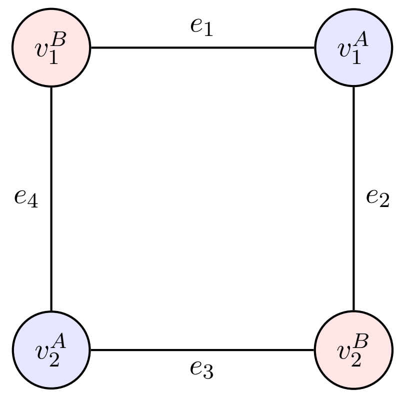

Figure 4. The swapped information graph G^τ associated with Figure 2.

subtrade on the cycle is mutually beneficial, whereas the subtrade off the cycle is not mutually beneficial, the component of the trade that creates mutual gains has signs that depend only on disagreement, and not on the assignment of information.

Example 2. Figure 4 differs from Figure 2 through the interchange of traders' information; once again, Trader A's information sets are blue and Trader B's information sets are red. As foreshadowed in the discussion of Example 1, Proposition 6 implies that the signs of any mutually agreeable trade in this new structure match those in the original structure. If the traders' priors remain $\mu^A = \left(\frac{3}{9}, \frac{1}{9}, \frac{2}{9}, \frac{3}{9}\right)$ and $\mu^B = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$ as in Example 1, then $\hat{t} = (3, -8, 7, -4)$ is a mutually agreeable trade in the swapped information graph in Figure 4. Note that $\hat{t}$ has the same signs as the mutually agreeable trade from Example 1, $t = (6, -7, 4, -5)$.

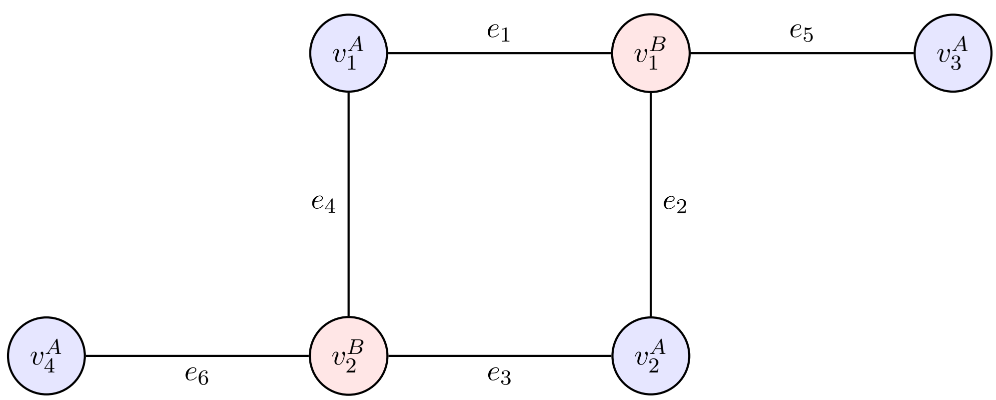

Figure 5. A single cycle with two additional edges and vertices.

Proposition 6 also describes directions of trade consistent with mutual agreeability when the information structure contains states outside of a unique cycle. In this case, the direction of trade off the cycle is *always* reversed. Our next example illustrates that part of Proposition 6, after which we present further definitions and results formalizing how the portion of trade that does flow in the direction of information does not derive from disagreement, nor does it create surplus for the two traders.

Figure 5 illustrates this case. It differs from Figure 2 through the addition of two edges, e_5, e_6, and two vertices, v_3^A, v_4^A. The unique cycle remains (e_1, e_2, e_3, e_4). Because A can discern e_5 and e_6, any mutually agreeable trade must have $t(e_5), t(e_6) > 0$. When information is swapped, any mutually agreeable trade must be negative on both of those states. On these states, trade *does* follow the direction of private information.

However, the mutually agreeable trade confined to those states off the cycle, namely confined to $\{e_5, e_6\}$, (coupled with 0 elsewhere) is not $\{e_5, e_6\}$-mutually agreeable. In the information structure of Figure 5, Trader B dislikes any trade that is positive on states e_5, e_6 and zero elsewhere.[22] In other words, the portion of trade that follows traders' information is *exactly* that portion of trade which traders do not find mutually agreeable.

Propositions 4 and 6 described cases (binary trades, and unicyclic information structures) where it is impossible to reverse all the signs of trade when interchanging traders' information. Nevertheless, it is possible in other settings for a trade to be mutually agreeable, and for a trade in the opposite direction to be mutually agreeable when information is swapped. However, we show below that this can only happen when trades are *inefficient* in the following sense.

Definition 6. *The trade t is* interim inefficient *if there exists some trade $\hat{t}$ such that (i) for all $e \in E$, $\text{sgn}(\hat{t}(e)) = \text{sgn}(t(e))$ and $|\hat{t}(e)| \leq |t(e)|$, and (ii) both parties weakly prefer $\hat{t}$ to t, and at least one trader strictly prefers $\hat{t}$ to t at some vertex.*[23]

In broad terms, we deem a trade to be interim inefficient whenever it is less desirable than a trade that shrinks the amounts transferred in some contingencies without increasing the amounts anywhere else.[24] We call a trade *interim efficient* if it is not interim inefficient.

We next establish two basic properties of interim-efficient trade. First, interim inefficiency is equivalent to a trade going against the direction of disagreement on some oriented cycle. Second, this efficiency criterion is not overly restrictive in the sense that every information graph that admits

[22]For a partial intuition for the general case, imagine removing all of the states in the cycle from the information structure. The ensuing acyclic information structure has no mutually agreeable trade for any beliefs, by Lemma 1 and Proposition 0.

[23]This definition uses the assumption that the graph is connected. When the graph is disconnected, the trade is inefficient if the condition holds on some connected component.

[24]In the context of an exchange economy, Wilson (1978) defines an allocation as lying outside of the *coarse core* if it is common knowledge that some feasible blocking trade is strictly preferred to the original trade. We impose a weaker criterion: it must be common knowledge that the blocking trade is weakly preferred to, but not indifferent to, the original trade. Otherwise, it would be impossible to rule out any trade that benefits a single trader across all states in one of their information sets, regardless of how inefficient that trade might appear elsewhere. For example, in Figure 5, for any mutually agreeable trade t, no weakly smaller trade can make A strictly better off in the singleton information set containing the state e_5. Our notion of inefficiency thus resembles lying outside of the "coarse core" of Wilson (1978) and the concept of "interim incentive inefficiency" introduced by Holmström and Myerson (1983).

a mutually agreeable trade admits a mutually agreeable trade that is interim efficient.

Proposition 7. *1. The trade t is interim inefficient if and only if it goes against the direction of disagreement on some oriented cycle.*

2. If there exists a mutually agreeable trade, then there exists a mutually agreeable and interim-efficient trade.

The trades described in Examples 1 and 2 are both interim efficient because they follow the direction of disagreement on the only oriented cycle.

Our next result expresses that if swapping information can flip the signs of mutually agreeable trade, then trade is interim inefficient under both the original and transposed information structures.

Proposition 8. *Let $G = (V, E, \partial_0, r, w)$ and $G' = (V', E, \partial_0', r', w)$ be information graphs with common edge set E, and suppose that their oriented versions $\vec{G} = (V, E, \partial_1, r, w)$ and $\vec{G}' = (V', E, \partial_1', r', w)$ have the same oriented cycles. If t and t' are mutually agreeable trades in G and G', respectively, and $\operatorname{sgn}(t(e)) = -\operatorname{sgn}(t'(e))$ for each $e \in E$, then both t and t' are interim inefficient.*

A special case of Proposition 8 involves swapping traders' information. This does not affect E and reverses the orientation of all cycles, thereby preserving the set of all oriented cycles.[25] By Proposition 2, every mutually agreeable trade must follow the direction of disagreement along some oriented cycle. Consequently, transposing traders' information changes every sign of trade only if each trade goes against the direction of disagreement along some oriented cycle, which is interim inefficient. Proposition 8 applies not just to swaps in information but to a broader class of transformations of information structures that preserve oriented cycles.

While our previous results show that swapping information cannot drastically change the direction of efficient trade—it cannot flip every sign of such a trade—we also address if and when trades with the *same* signs can be mutually agreeable and interim efficient under the original and

[25]For example, Figures 2 and 4 depict information structures that produce the same set of oriented cycles; the edges e_1 and e_3 share signs in $W(\cdot)$ under either orientation of the single cycle. The information structure below gives rise to a different set of oriented cycles since e_1 and e_3 no longer share signs in $W(\cdot)$:

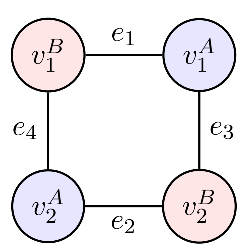

transposed information structures. This is indeed possible, and our next results provide conditions for when swapping information preserves every sign of efficient trade.

To state our next result, we use the standard graph-theoretic terminology that a graph is *bridgeless* if it contains no edge whose removal increases the number of connected components of the graph; equivalently, the graph is bridgeless if every edge of the graph belongs to some cycle.

Lemma 3. *Let G be a bridgeless graph that contains an oriented cycle $\vec{C}$ for which $W(\vec{C}) \neq 0$. Then there exist trades t and t^τ that are mutually agreeable and interim efficient in G and G^τ, respectively, and share the same sign pattern.*

The proof of Lemma 3 works by representing the signs of a trade through an orientation of the graph, and showing that there exists an orientation for which every vertex is reachable from every other vertex by a directed path. This ensures that the surplus generated by trade on the cycle with disagreement can be propagated through the graph, producing mutual agreeability. The sufficient condition in Lemma 3 matches the necessary and sufficient condition of Robbins (1939) for the existence of a strongly connected orientation of an undirected graph (under our maintained assumption that the graph is connected). It is satisfied in the cyclical graphs of Examples 1 and 2, and also the information structures assumed by Hellwig (2013) in the two-trader case. The bulk of the proof consists of showing that the mutually agreeable trades whose signs correspond to the orientation can be made interim efficient. When information is swapped, only the magnitudes—and not the signs—of trades must change.

In the example of Figure 5, states e_5 and e_6 do not lie on any cycle, so the sufficient condition for Lemma 3 does not hold. In this setting, the signs of mutually agreeable trade cannot be preserved when information is swapped. In the original structure, $t(e_6) > 0$ is necessary for A to accept the trade at v_4^A, whereas $t(e_6) < 0$ is necessary for B to accept the trade at that vertex. However, as discussed after Proposition 6, *some* of the signs of trade will be preserved when information is swapped in this example. Building on this intuition, we can extend the result from Lemma 3 to more general information graphs when we focus on portions of trade that generate mutual gains by themselves.

To develop this result, first notice that every mutually agreeable trade t has some set $S \subseteq E$ such that its restriction t_S is mutually agreeable on the edge-induced subgraph G_S, since $S = E$ always works. In settings like that in Example 1, mutually agreeable trade requires trading in every state, and $S = E$ is a minimal support of any mutually agreeable trade. In other settings, like that of Figure 5, by contrast, the restriction of a mutually agreeable trade to $\{e_1, e_2, e_3, e_4\}$ can itself be mutually agreeable on the corresponding edge-induced subgraph. The original trade is therefore also $\{e_1, e_2, e_3, e_4\}$-mutually agreeable, and the states e_5 and e_6 contribute nothing to this mutual agreeability; essentially, for trades on G, they drain surplus from the cycle. Thus, $\{e_1, e_2, e_3, e_4\}$

is the only minimal-mutually-agreeable support in the example, and it is a cycle. This observation motivates our next definition.

Definition 7. *A set $S \subseteq E$ is a* minimal-mutually-agreeable support *if there exists a trade t_S such that (i) t_S is S-mutually agreeable; and (ii) for every nonempty $S' \subsetneq S$, the restriction $t_{S'}$ is not S'-mutually agreeable.*

A minimal-mutually-agreeable support identifies a set of states on which a trade can generate mutual gains by itself, but for which no proper nonempty restriction of that trade can do so.[26]

Although all signs of the trade cannot necessarily be preserved under an information swap, the signs of the surplus-generating portion of the trade can be maintained. In other words, the signs of the surplus-generating portion of trade follow disagreement rather than information. Proposition 9 formalizes this observation.

Proposition 9. *For each minimal-mutually-agreeable support S, there exist trades t and t^τ that are mutually agreeable and interim efficient in G and G^τ, respectively, and share the same sign pattern on S.*

A special case of the bridgeless condition (i.e., every edge belongs to a cycle) occurs whenever traders face residual uncertainty over which they disagree after joining their information. In this case, we obtain a stronger result: there exists an interim-efficient mutually agreeable trade that remains mutually agreeable when the traders' information is swapped.

Proposition 10. *If every pair $(v^A, v^B) \in V^A \times V^B$ has non-constant edge weights on the edge-induced subgraph $G_{E(v^A)\cap E(v^B)}$, whenever $E(v^A) \cap E(v^B) \neq \emptyset$, then there exists a trade t that is mutually agreeable and interim efficient under G and G^τ.*[27]

Next, we offer another sufficient condition for the same trade to be mutually agreeable and interim efficient under an original and swapped information structure. This condition describes a form of uniformity in the nature of the traders' disagreement along a cycle. We deem the traders' disagreement to be *monotone* on a cycle whenever it follows the same direction on every information set within the cycle as on the entire cycle.

[26]Note that not every trade that is mutually agreeable on G_S makes S a minimal-mutually-agreeable support. The following example illustrates how minimal-mutually-agreeable supports, because they may be associated with different trades, can differ and even be ordered by set inclusion. Let $E = \{e_1, e_2, e_3\}$, and suppose that the two traders have priors $\mu^A = \left(\frac{9}{20}, \frac{4}{20}, \frac{7}{20}\right)$ and $\mu^B = \left(\frac{5}{20}, \frac{6}{20}, \frac{9}{20}\right)$. The trade $t = (4, -1, -2)$ is mutually agreeable, yet no two-state restriction of t is mutually agreeable: E is a minimal-mutually-agreeable support. Now consider $t' = (3, -1, -2)$, which is also mutually agreeable. Unlike t, the trade t' remains mutually agreeable when restricted to the two-state support $\{e_1, e_3\}$. Hence, both $S = \{e_1, e_2, e_3\}$ and $S' = \{e_1, e_3\}$ are minimal-mutually-agreeable supports (albeit associated with different trades).

[27]Formally, for each $v^A, v^B \in V$ such that $E(v^A) \cap E(v^B) \neq \emptyset$, $\{w(e) : e \in E(v^A) \cap E(v^B)\}$ has at least two elements; namely, beliefs are not collinear on any element of the join of the traders' information partitions.

Definition 8. *Let $\vec{C}$ be an oriented cycle. Disagreement is* monotone *on $\vec{C}$ if, for every pair of adjacent edges $e, e' \in C$,*

$$W(\vec{C}) > 0 \quad \Rightarrow \quad \sigma_{\vec{C}}(e)w(e) + \sigma_{\vec{C}}(e')w(e') > 0,$$

and

$$W(\vec{C}) < 0 \quad \Rightarrow \quad \sigma_{\vec{C}}(e)w(e) + \sigma_{\vec{C}}(e')w(e') < 0.$$

Disagreement along a cycle consisting of two edges is always monotone. Longer cycles exhibit monotone disagreement if the direction of disagreement along the entire cycle holds along every pair of adjacent edges. For example, if $W(\vec{C}) > 0$—the geometric average of A's likelihood ratios of cycle edges oriented towards A vertices relative to those oriented towards B vertices exceeds the corresponding geometric average for B—then monotonicity requires that for any pair of adjacent edges, A has a higher likelihood ratio of the edge oriented towards the A vertex relative to the edge oriented towards the B vertex than does B. The traders' beliefs in Example 1 are not monotone because $W(\vec{C}) < 0$, yet A assigns a higher odds ratio of e_4, which is traversed towards an A vertex, to e_3, which is traversed towards a B vertex, than does B.

Proposition 11. *Suppose disagreement is monotone on some oriented cycle $\vec{C}$ with $W(\vec{C}) \neq 0$. Then there exists a trade t that is C-mutually agreeable and interim efficient under both G and G^τ.*

To illustrate the result, let us return to Example 1 and modify A's beliefs to become $\mu^A = \left(\frac{3}{7}, \frac{1}{7}, \frac{2}{7}, \frac{1}{7}\right)$, while keeping Trader B's priors the same at $\mu^B = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$. The disagreement has become monotone, and Proposition 11 implies the existence of a trade that is mutually agreeable under original and interchanged information structures, namely those depicted in Figures 2 and 4. The trade $t = (4, -5, 4, -5)$ is one such trade. More generally, when the information structure contains information sets beyond those in the cycle C with monotone disagreement, we can find a trade that is C-mutually agreeable under the original and interchanged information structures. For instance, in Figure 5, if traders' beliefs conditional on being in states 1 to 4 are as above, then disagreement on the cycle is monotone. Hence, for sufficiently small values of $w, x, y, z > 0$, the trade $t = (4, -5, 4, -5, w, x)$ is mutually agreeable under the original information structure, and $t' = (4, -5, 4, -5, -y, -z)$ is mutually agreeable under the interchanged information structure.

4 Extensions

This section discusses how the principles established above extend to two alternative environments.

4.1 Multiple Traders

Our analysis above outlines how mutually agreeable trades depend on disagreements and information structures, while abstracting from the institutional details of markets that give rise to more specific predictions. Although our motivating contexts involve retail investors in large markets, one key abstraction has been to consider only two traders. Although we believe that we could extend our analysis to more complex environments with many traders, we limit our analysis here to showing how the qualitative features of Proposition 5 carry through.[28]

To consider this question, we define multilateral trade by imagining that each of many traders faces a trade that determines their state-contingent payoff and ask when such trades are mutually agreeable. More formally, for a set of traders $i = 1, 2, \ldots, I$, let $t^i : E \to \mathbb{R}$, where $t^i(e)$ denotes the transfer to i ($t^i(e) \geq 0$) or away from i ($t^i(e) \leq 0$) in each state e, where $\sum_{i=1}^{I} t^i(e) = 0$.

We first address an analog to Proposition 5 in the multilateral setting: starting from a mutually agreeable trade, can we swap the information partitions and trades of two traders while preserving mutual agreeability?

Example 3. Traders A, B, C bet on whether a stock will move significantly upward, stay roughly the same, or move significantly downward. That is, $E = \{\text{Up}, \text{Same}, \text{Down}\}$. A and B both perceive a 50% chance of a large price increase and a 49% chance of no significant change, whereas C foresees the price dropping with 99% probability (and going up or staying the same with equal probability). Trader C bets against both A and B, at even odds and for $\$m$, that the price will drop. In addition, A and B make a side bet in which A pays $q < m$ to B if the price increases, while B pays q to A if the price stays roughly the same. More formally, $\mu^A = \mu^B = (0.50, 0.49, 0.01)$ and $\mu^C = (0.005, 0.005, 0.99)$; $t^A = (m - q, m + q, -m), t^B = (m + q, m - q, -m)$ and $t^C = (-2m, -2m, 2m)$. Trade is mutually agreeable. Traders A and B also prefer *each other's* positions to no trade, since each expects to gain enough from C to offset any loss on the side bet. Therefore, exchanging their (trivial) information partitions does not change their preference for trade.

Example 3 illustrates the possibility of non-trivially permuting traders' information partitions and trades while still preserving mutual agreeability. However, a variant of Proposition 5 does hold with an additional assumption that is highly relevant for real-world trading. Consider the same three-trader scenario, but now suppose that C is an institutional trader, and A and B are small-scale (retail) traders, who agree that C—"smart money"—does not lose money on average from trading.[29]

To present our negative result in this three-trader setting, it suffices to consider two traders, A and B, and use the same information graph as before. If Traders A and B do not find a trade

[28]We can formalize the information structure by generalizing our multigraph to an n-uniform, n-partite hypergraph: each edge is incident to n vertices, one from each of the n elements of the partition of V corresponding to the n traders.

[29]In some formal market models, C could be a market-maker.

mutually agreeable, then adding Trader C cannot make the trade mutually agreeable for all three traders.

Corollary 2. *Assume that for each $i \in \{A, B\}$ and every $v \in V$, $\mathbb{E}^i\left[-(t^A(e) + t^B(e))|e \in E(v)\right] > 0$. If $t = (t^A, t^B)$ is mutually agreeable in G, then $t' = (t^B, t^A) + (k\mathbf{1}, -k\mathbf{1})$ is not mutually agreeable in G^τ for any $k \in \mathbb{R}$.*

Note that the logic of Corollary 2 extends straightforwardly to cases with more than two small traders, provided that each pair of traders believes that their joint trade generates revenue for the institutional trader.

Moreover, Corollary 2 has implications for insurance markets in settings with small-scale risk. Customers may disagree about their own and others' likelihood of loss, but they likely agree that the insurance company does not lose money on average, and so the same logic applies.

4.2 Non-Zero-Sum Common-Value Trade

Although we have focused on zero-sum trades, the basic idea that trade should not follow the direction of information extends to some positive-value settings. Of course, most surplus-generating trade in goods and services doesn't come from different beliefs, but from different tastes, needs, and ability to pay. But we believe the logic we have spelled out in zero-sum settings may be quite relevant in many contexts where people have very similar and obvious goals. For instance, what explains why different consumers make such different choices regarding health practices and medical treatments when they all want to cure what ails them? When people more or less share a common taste for some good or service as a function of beliefs about their effectiveness, how might we explain different consumers implementing such different beliefs?

To illustrate how one can deploy the approach of this paper to address such behavior, consider two patients who each decide whether to use a medical treatment with uncertain but common value. Let the common value of the treatment be $v(e)$ in state e, and patient i's utility without the treatment is $u^i(e)$. A patient's payoff from using the treatment is then $v(e) + u^i(e)$. Suppose the two patients have different priors over the state—they disagree on the expected value of the treatment—and the information structure is represented by information graph G. An immediate corollary of Proposition 5 applies to this setting.

Corollary 3. *Suppose that under the information graph G with edge set E, Patient A prefers treatment to no treatment at each $v^A \in V^A$, whereas Patient B prefers no treatment to treatment at each $v^B \in V^B$. Then under any connected information graph $\widehat{G}$ that has the same edge set E and edge weights w, it cannot be the case that Patient B prefers treatment to no treatment at each $\hat{v}^B \in \widehat{V}^B$, whereas Patient A prefers no treatment to treatment at each $\hat{v}^A \in \widehat{V}^A$.*

The intuition mirrors the logic of Proposition 5. To frame this positive-value trade in our previous notation, define $\tilde{u}^i(e) := u^i(e) + v(e)/2$ and let $t(e) = v(e)/2$. When A accepts the treatment but B declines, then A earns $t(e)$ and B earns $-t(e)$.[30] In this way, disagreement over the efficacy of a medical treatment formally resembles zero-sum speculation.

5 Discussion and Conclusion

This paper examines a simple model with heterogeneous priors to clarify what such models do and don't predict. While these models potentially offer an explanation for a large volume of speculative trade, they also impose restrictions on how agreeable trades depend on traders' information. In particular, one thing these models do not predict is people systematically trading in the direction of the news they happen to receive—e.g., betting on a firm after a favorable report and against it after a poor one—which is an important feature of behavior among small-scale traders (Hong and Stein, 1999; Malmendier and Shanthikumar, 2007). More broadly, we establish a series of results showing the limited extent to which trade from a heterogeneous-priors perspective should follow the direction of private information. These findings imply that if information is distributed among traders independently of their priors, we should not expect much correlation between the content of a trader's information and the direction they trade.[31] Such a correlation can arise, however, by relaxing rational inference about what others' willingness to trade reveals. Our supplementary analysis in Appendix A shows that failures of inference (as in Eyster and Rabin, 2005) can explain trading patterns that heterogeneous priors cannot.[32] The findings offer a useful contrast with heterogeneous-priors-based trade.[33]

Throughout the paper, we have assumed that the traders can make any trade $t : E \to \mathbb{R}$. This assumption may lack external validity. For example, it may not be feasible to condition trade on the traders' signals. Such restrictions could be incorporated into our framework by assuming that

[30]Strictly speaking, the antecedents of Proposition 5 and Corollary 3 differ by virtue of the former assuming that A prefers $t(e)$ to no trade and the latter assuming that A prefers $t(e)$ to $-t(e)$. However, these are equivalent given the linearity of payoffs.

[31]There are cases where one might argue that common-knowledge disagreements could be correlated with who observes information. For instance, some intuitions about the meaning of "overconfidence" could be interpreted as people believing that ideas that pop into their own heads are wiser than other people think they are. It would be hard to make sense of swapping information in such cases, and our framework has limited applicability to trade driven by this form of private information.

[32]Our analysis in the appendix can be thought of as complementary to the analysis in Eyster et al. (2019), who argue that important patterns in financial markets (most notably, the volume of trade) are better explained by the neglect of others' beliefs implicit in market prices rather than disagreement with those beliefs.

[33]On the other hand, in part because the framework of this paper is ill-suited to the study of cursedness, the set of outcomes ruled out by our cursed criterion is fairly small (and it never rules out mutually agreeable trade). Thus, while our primary focus is on the comparative statics with respect to information, our general ability to accommodate trade that heterogeneous priors cannot should be taken with a grain of salt.

t must be measurable with respect to some partition $\mathscr{P}$ of E. Some of our results clearly depend upon the absence of contracting restrictions. For example, Lemma 2 requires a sufficiently rich contracting space. Trivially, if $\mathscr{P}$ consists of a single set-element—$t(e) = c \in \mathbb{R}$ for all $e \in E$—then no mutually agreeable trade is possible. However, some of our main results—demonstrating that gainful trade does not follow traders' information—do not require a rich contracting space.

We have additionally assumed throughout that the traders are risk neutral. Risk-averse traders who have heterogeneous priors may also find speculative trade mutually agreeable. Formally, consider two risk-averse traders with endowments and Bernoulli utility functions that do not vary with the state, as well as their risk-neutral counterparts who share their respective priors. A risk-averse trader who strictly prefers t to no trade assigns t a strictly positive certainty equivalent, which does not exceed the expected value of the trade. Their risk-neutral counterpart therefore prefers t to no trade.

Consequently, many of our results carry over to risk-averse traders. For example, the conclusion of Proposition 2 that mutually agreeable trade must follow the direction of disagreement along some cycle holds for risk-averse traders. Since risk aversion preserves the preference for the trade contractions underlying our notion of interim efficiency, our central result also continues to hold: if swapping information reverses the signs of trade, then those trades are interim inefficient (Proposition 8).

References

CARRILLO, J. AND T. PALFREY (2011): "No Trade," *Games and Economic Behavior*, 71, 66–87.

CERREIA-VIOGLIO, S., R. CORRAO, AND G. LANZANI (2025): "(Un-)Common Preferences, Ambiguity, and Coordination," Mimeo.

COHEN, S. AND S. LI (2026): "Sequential Cursed Equilibrium," *American Economic Review*, 116, 934–976.

CRAWFORD, V. P. AND N. IRIBERRI (2007): "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica*, 75, 1721–1770.

DOW, J., V. MADRIGAL, AND S. WERLANG (1990): "Preferences, Common Knowledge and Speculative Trade," Working paper.

EYSTER, E. AND M. RABIN (2005): "Cursed Equilibrium," *Econometrica*, 73, 1623–1672.

EYSTER, E., M. RABIN, AND D. VAYANOS (2019): "Financial Markets Where Traders Neglect the Informational Content of Prices," *The Journal of Finance*, 74, 371–399.

FEINBERG, Y. (1995): "A Converse to the Agreement Theorem," Working Paper 83, Center for Rationality and Interactive Decision Theory, The Hebrew University, Jerusalem.

FONG, M.-J., P.-H. LIN, AND T. R. PALFREY (2025): "Cursed Sequential Equilibrium," *American Economic Review*, 115, 2616–2658.

GEANAKOPLOS, J. (1994): "Common knowledge," in *Handbook of Game Theory with Economic Applications*, ed. by R. Aumann and S. Hart, Elsevier, vol. 2, chap. 40, 1437–1496.

GIZATULINA, A. AND Z. HELLMAN (2019): "No trade and yes trade theorems for heterogeneous priors," *Journal of Economic Theory*, 182, 161–184.

GLEISS, P., J. LEYDOLD, AND P. F. STADLER (2003): "Circuit Bases of Strongly Connected Digraphs," *Discussiones Mathematicae Graph Theory*, 23, 241–260.

GOLUB, B. AND S. MORRIS (2017): "Higher-Order Expectations," Mimeo.

HARSANYI, J. C. (1968): "Games with Incomplete Information Played by 'Bayesian' Players, I-III. Part II. Bayesian Equilibrium Points," *Management Science*, 14, 320–334.

HELLMAN, Z. AND D. SAMET (2012): "How Common are Common Priors?" *Games and Economic Behavior*, 74, 517–525.

HELLWIG, M. (2013): "From posteriors to priors via cycles: An addendum," *Economics Letters*, 118, 455–458.

HOLMSTRÖM, B. AND R. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.

HONG, H. AND J. C. STEIN (1999): "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *The Journal of Finance*, 54, 2143–2184.

JEHIEL, P. AND F. KOESSLER (2008): "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *Games and Economic Behavior*, 62, 533–557.

MALMENDIER, U. AND D. SHANTHIKUMAR (2007): "Are Small Investors Naive About Incentives?" *Journal of Financial Economics*, 85, 457–489.

MILGROM, P. AND N. STOKEY (1982): "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26, 17–27.

MORRIS, S. (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327–1347.

MORRIS, S. AND C. SKIADAS (2000): "Rationalizable Trade," *Games and Economic Behavior*, 31, 311–323.

NYARKO, Y. (2010): "Most games violate the common priors doctrine," *International Journal of Economic Theory*, 6, 189–194.

ROBBINS, H. (1939): "A Theorem on Graphs, with an Application to a Problem of Traffic Control," *The American Mathematical Monthly*, 46, 281–283.

RODRIGUES-NETO, J. A. (2009): "From Posteriors to Priors via Cycles," *Journal of Economic Theory*, 144, 876–883.

——— (2012): "The Cycles Approach," *Journal of Mathematical Economics*, 48, 207–211.

SAMET, D. (1998a): "Common Priors and Separation of Convex Sets," *Games and Economic Behavior*, 24, 172–174.

——— (1998b): "Iterated Expectations and Common Priors," *Games and Economic Behavior*, 24, 131–141.

SEBENIUS, J. K. AND J. GEANAKOPLOS (1983): "Don't Bet on it: Contingent Agreements with Asymmetric Information," *Journal of the American Statistical Association*, 78, 424–426.

TIROLE, J. (1982): "On the Possibility of Speculation under Rational Expectations," *Econometrica*, 50, 1163–1181.

VEBLEN, O. (1912): "An Application of Modular Equations in Analysis Situs," *Annals of Mathematics*, 14, 86–94.

WILSON, R. (1978): "Information, Efficiency, and the Core of an Economy," *Econometrica*, 46, 807–816.

A Cursed Trade

In this appendix, we outline an approach to investigating "cursedness" (Eyster and Rabin, 2005)—where traders fail to account for others' private information when inferring from their willingness to trade—in the bilateral trading setting considered throughout the paper.[34] We characterize the implications of cursedness and contrast those implications to the heterogeneous-priors approach above. In particular, unlike trade driven by heterogeneous priors, we show that the direction of trade fueled by cursedness can follow the direction of private information but is largely unrelated to disagreements in priors.

Consider again the example in Panel (a) of Figure 1, reproduced here:

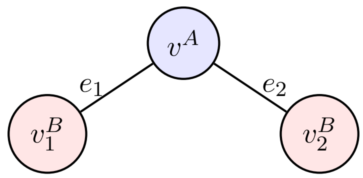

Since the graph is acyclic, no mutually agreeable trade exists, irrespective of priors. Trader A will not agree to any trade where $t(e) < 0$ in either state, so trade that favors B in either state is not plausible. However, if A is "cursed", then trades with $t(e) < 0$ become plausible.

Consider a trade under which $t(e_1) > 0 > t(e_2)$, together with a strategy for B that involves accepting the trade in e_2 and rejecting in e_1. A would be willing to trade as long as they believe that the likelihood of e_1 is sufficiently high conditional on trade occurring. In particular, a "cursed" Trader A, who we define as neglecting the relationship between B's willingness to trade and B's information set (the state), would be willing to trade given a sufficiently high prior on e_1.

For $e \in E$, define $v^j(e)$ to be the unique vertex in $\partial_0(e)$ with $r(v^j(e)) = j$, and let $E^j(e) := E(v^j(e))$. $E^j(e)$ is Trader j's information set in state e. Additionally, let

$$S^A(t) := \{e \in E : \mathbb{E}_{\mu^A}[t(e')|e' \in E^A(e)] > 0\};$$

this is the set of states in which Trader A regards the trade as strictly preferable to no trade, based on their information alone. Likewise, let

$$S^B(t) := \{e \in E : \mathbb{E}_{\mu^B}[t(e')|e' \in E^B(e)] < 0\}$$

be the set of states in which Trader B regards the trade as strictly preferable to no trade, based on their information alone.

Definition A.1. *The trade t is* cursed rationalizable *in state $e \in E$ if $e \in S^A(t) \cap S^B(t)$.*

A trade is cursed rationalizable in state e if both traders find the trade profitable at e when they believe the other party will also agree to the trade. This state-by-state definition of plausible trade

[34]Other models that capture failures of inference from others' behavior, such as analogy-based expectations (Jehiel and Koessler, 2008) and level-1 thinking (Crawford and Iriberri, 2007), lead to conclusions similar to those we discuss here.

differs from mutual agreeability by not requiring that both traders agree to the trade everywhere. A mutually agreeable trade is cursed rationalizable in all states.[35]

We are interested in comparing the directions of cursed trade and mutually agreeable trade.

Definition A.2. Trade that favors A is cursed rationalizable in state $e \in E$ *if there exists some trade* t *that is cursed rationalizable in state* e*, where* $t(e) > 0$*.* Trade that favors B is cursed rationalizable in state $e \in E$ *if there exists some trade* t *that is cursed rationalizable in state* e*, where* $t(e) < 0$*.*

The order of quantifiers permits trade favoring A and favoring B in the same state.

Proposition A.1.

1. *If* $E^A(e) = E^B(e)$*, then trade that favors* i *at* e *is cursed rationalizable if and only if there exists a mutually agreeable trade in* $G_{E^A(e)}$ *that favors* i *at* e*.*
2. *If* $E^i(e) = \{e\}$*, then trade that favors* $j \neq i$ *is not cursed rationalizable at* e*.*
3. *If* $\{e\} \neq E^i(e) \neq E^j(e)$*, then trade that favors* j *is cursed rationalizable at* e*.*

If the two traders have the same information at e, then $E^A(e) = E^B(e)$. This implies that $G_{E^A(e)}$ is a connected component of the information graph G. If both like t at some state in $E^i(e) = E^j(e) =: E(e)$, then both must like t at all states in $E(e)$, and hence t must be mutually agreeable. Additionally, like in the example above, if Trader i learns e, then they cannot lose in any cursed rationalizable trade at e.

Our next result shows that whenever traders have different information at a state, the feasible directions of cursed trade there are independent of the traders' prior beliefs.

Corollary A.1. *Suppose* $E^A(e) \neq E^B(e)$*. Consider any two pairs of prior beliefs* (μ^A, μ^B) *and* $(\tilde{\mu}^A, \tilde{\mu}^B)$*, and let* $w = \ln \mu^A - \ln \mu^B$ *and* $\tilde{w} = \ln \tilde{\mu}^A - \ln \tilde{\mu}^B$*. Trade that favors Trader* i *in state* e *is cursed rationalizable under* (μ^A, μ^B) *on the information graph* $G = (V, E, \partial_0, r, w)$ *if and only if trade that favors Trader* i *in state* e *is cursed rationalizable on* $(\tilde{\mu}^A, \tilde{\mu}^B)$ *on* $\tilde{G} = (V, E, \partial_0, r, \tilde{w})$*.*

In terms of our graphical representation, changing the disagreement weights on the information graph while holding its information structure fixed does not change which directions of trade are cursed rationalizable whenever the traders have different information.

In Section 2 of the main text, we showed that giving traders additional information can undermine trade under heterogeneous priors (Proposition 3), just as it does under homogeneous priors.

[35] Alternatively, extending the notion of fully cursed equilibrium of Eyster and Rabin (2005) to encompass heterogeneous priors, a trade t is cursed rationalizable in state e if, in some fully cursed equilibrium, both traders accept t in state e. In this trading game, both traders have action spaces $\{\text{Accept}, \text{Reject}\}$, and trade occurs in e if and only if both traders choose Accept in e; otherwise, payoffs are 0. Although the definition of cursed equilibrium assumes that each party understands that the other party would only sometimes agree to the trade, treating the other's agreement as independent of their information leads to the same optimization as believing they always agree to the trade.

Refining traders' information works very differently in cursed-rationalizable trade. An intuition for cursedness in general is that people are better off if they have more private information when embedded in an economy where others aren't paying attention to what the trading behavior of others reveals about these others' private information. While having private information is not necessary for the possibility that a trader can gain in a particular situation, it is sufficient. Our final result, stemming from Part 3 of Proposition A.1, demonstrates how private information influences the direction of cursed trades. We say that Trader i has private information if revealing their information makes Trader j more informed.

Corollary A.2. *If i has private information at e, then trade that favors i is cursed rationalizable in state e.*

In contrast to our results under heterogeneous priors, this result emphasizes that there is always cursed-rationalizable trade that flows in the direction of a trader's private information.

B Graphs and Partitions

In this section, we draw the four information structures that appear in figures of the main text as both graphs and partitions in order to illustrate the two approaches side by side.

Figure B.1 depicts the two-state example in Figure 1. Figure B.2 depicts the 4-state cyclic setting of Figure 2. Figure B.3 depicts the 4-state acyclic setting of Figure 3. Figure B.4 depicts the 6-state setting of Figure 5 that has one cycle of length four.

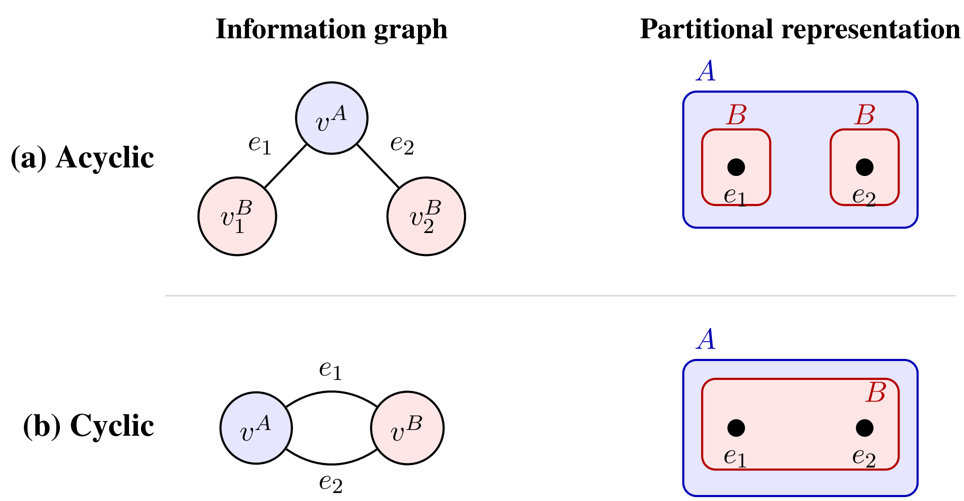

Figure B.1. A two-state information structure in which Trader A has a single information set. In panel (a), A cannot distinguish the two states, whereas B can; the information graph is acyclic. In panel (b), neither agent can distinguish the two states; there is a cycle of length 2.

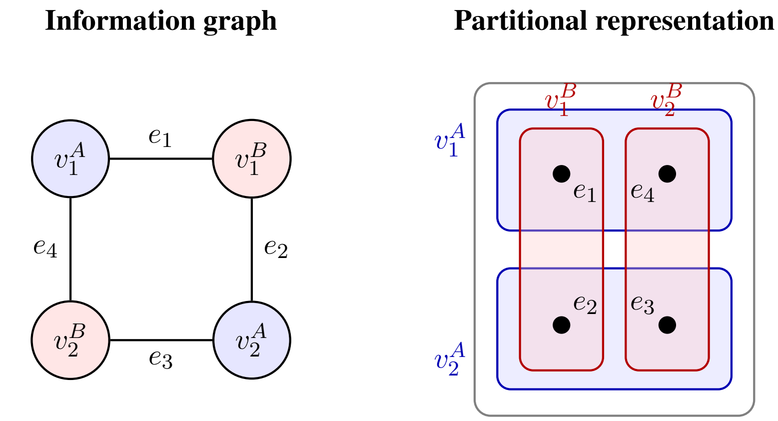

Figure B.2. A four-state information structure forming a cycle. Trader A's information partition is $\{\{e_1, e_4\}, \{e_2, e_3\}\}$, while Trader B's information partition is $\{\{e_1, e_2\}, \{e_3, e_4\}\}$.

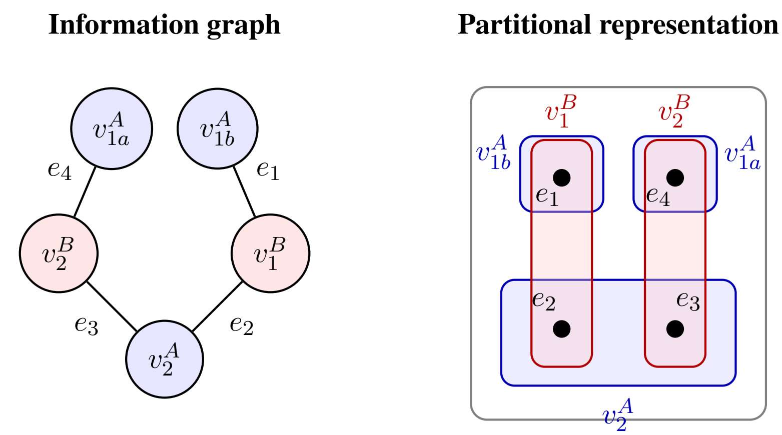

Figure B.3. Trader A's information partition is $\{\{e_1\}, \{e_4\}, \{e_2, e_3\}\}$, and Trader B's information partition is $\{\{e_1, e_2\}, \{e_3, e_4\}\}$.

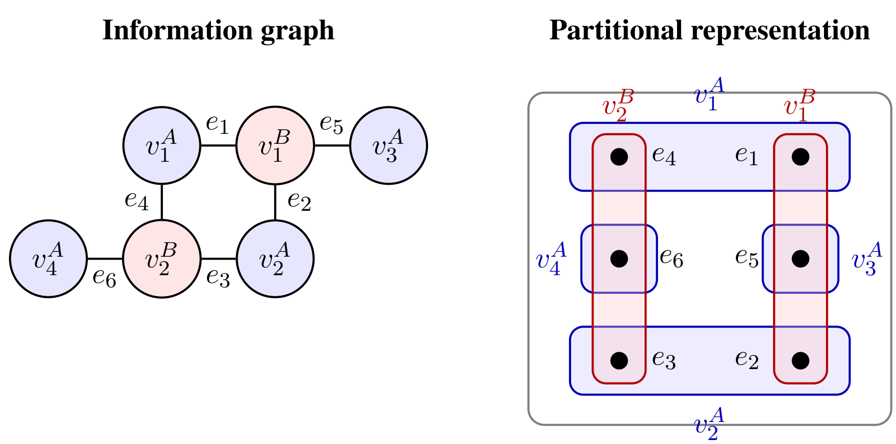

Figure B.4. Trader A's information partition is $\{\{e_1, e_4\}, \{e_2, e_3\}, \{e_5\}, \{e_6\}\}$, while Trader B's information partition is $\{\{e_1, e_2, e_5\}, \{e_3, e_4, e_6\}\}$.

C Proofs

Proof of Proposition 0. Although the first statement in Proposition 0 was already proven by Rodrigues-Neto (2009, 2012) and Hellman and Samet (2012), and the second statement was proven by Nyarko (2010) and Hellman and Samet (2012), we provide a shorter and different proof in our setting that makes use of the information graph.

Given the information graph G and its reference orientation, define its $|V| \times |E|$ incidence matrix B, whose ve-th element is $+1$ if e points into v, -1 if e points out of v and is 0 otherwise. The kernel of B,

$$\ker(B) = \{w \in \mathbb{R}^{|E|} : Bw = 0\},$$

is the *flow space* of the oriented graph. The image of B transpose,

$$\mathrm{Im}(B^\top) = \{w \in \mathbb{R}^{|E|} : w = B^\top x \text{ for some } x \in \mathbb{R}^{|V|}\}$$

is its *cut space*. The vector x can be interpreted as assigning *potentials* to each vertex $v \in V$. Since the row in $B^\top$ corresponding to the edge e, with $\partial_1(e) = v - u$, has -1 in the coordinate corresponding to u, $+1$ in the coordinate corresponding to v, and zeros in all other coordinates,

$$w(e) = x(v) - x(u).$$

Thus $w \in \mathrm{Im}(B^\top)$ if and only if all of its edge weights can be expressed as differences in vertex potentials.

We claim that this is equivalent to the existence of a common prior. For $\partial_1(e) = v - u$,

$$\begin{aligned} \ln \mu^A(e) - \ln \mu^B(e) &= x(v) - x(u) \\ \Rightarrow \mu^A(e) e^{-x(v)} &= \mu^B(e) e^{-x(u)} =: q(e) \end{aligned}$$

Normalize these positive weights to obtain a probability distribution

$$\mu(e) := \frac{q(e)}{\sum_{e' \in E} q(e')}.$$

We will show that μ is a common prior consistent with both traders' interim beliefs. To see this, fix an A-vertex v. Since the factor $e^{-x(v)}$ is constant across all $e \in E(v)$,

$$\mu(e|E(v)) = \frac{q(e)}{\sum_{e' \in E(v)} q(e')} = \frac{\mu^A(e) e^{-x(v)}}{\sum_{e' \in E(v)} \mu^A(e') e^{-x(v)}} = \frac{\mu^A(e)}{\mu^A(E(v))} = \mu^A_{E(v)}(e).$$

Likewise, for any B-vertex u, the factor $e^{-x(u)}$ is constant across all $e \in E(u)$, and hence

$$\mu(e|E(u)) = \frac{\mu^B(e)}{\mu^B(E(u))} = \mu^B_{E(u)}(e).$$

Thus, μ induces both traders' interim beliefs and is therefore a common prior consistent with those beliefs.

Conversely, suppose that there exists a common prior μ consistent with both traders' interim beliefs. For every A-vertex v, there exists some $\alpha_v > 0$ such that $\mu(e) = \alpha_v \mu^A(e)$ for all $e \in E(v)$, and for every B-vertex u, there exists some $\beta_u > 0$ such that $\mu(e) = \beta_u \mu^B(e)$ for all $e \in E(u)$. Thus, for any edge e with $\partial_1(e) = v - u$, $\mu^A(e)/\mu^B(e) = \beta_u/\alpha_v$, and therefore $w(e) = \ln \mu^A(e) - \ln \mu^B(e) = \ln \beta_u - \ln \alpha_v$. Setting $x(v) := -\ln \alpha_v$ for $v \in V^A$ and $x(u) := -\ln \beta_u$ for $u \in V^B$ gives $w(e) = x(v) - x(u)$. Hence $w = B^\top x$, so $w \in \mathrm{Im}(B^\top)$.

By the Fundamental Theorem of Linear Algebra, $(\ker(B))^\perp = \mathrm{Im}(B^\top)$. Hence,

$$\dim(\mathrm{Im}(B^\top)) + \dim(\ker(B)) = |E|.$$

If G is acyclic, then $\ker(B) = \{0\}$, so $\mathrm{Im}(B^\top) = \mathbb{R}^{|E|}$. Thus, every vector of edge disagreements w belongs to $\mathrm{Im}(B^\top)$ and, by the equivalence established above, the traders' interim beliefs are consistent with a common prior. This proves Part 1.

Now suppose G is cyclic. Then $\ker(B) \neq \{0\}$, so $\dim(\mathrm{Im}(B^\top)) < |E|$. Hence the cut space $\mathrm{Im}(B^\top)$ is a proper linear subspace of $\mathbb{R}^{|E|}$ and has Lebesgue measure 0. To translate this conclusion into a statement about prior pairs, fix any oriented cycle $\vec{C}$ in G. Consistency with a common prior requires $W(\vec{C}) = 0$. Viewed as a function of (μ^A, μ^B) on the interior of $\Delta(E) \times \Delta(E)$, $W(\vec{C})$ is a non-constant real analytic function. Its zero set therefore has Lebesgue measure zero. Since the set of prior pairs whose interim beliefs are consistent with a common prior is contained in this zero set, that set also has Lebesgue measure zero. ■

Proof of Lemma 1. First, suppose that G is acyclic. Because G is connected, it is a tree, so $|E| = |V| - 1$. For any edge e, at most one of its endpoints can receive a positive payoff from trade. Hence, at most $|E| = |V| - 1$ vertices can receive a positive payoff from an incident edge. However, Conditions 1 and 2 of Lemma 1 require such an incident edge for every vertex in V, a contradiction.

Conversely, suppose that G contains a cycle C, and fix an orientation σ_C of C as in Definition 2. For each $e \in C$, set $t(e) = \sigma_C(e)$. At each A-vertex on C, one of its two incident cycle edges is traversed in the reference direction, and hence has $\sigma_C(e) = +1$, so that $t(e) > 0$. At each B-vertex on C, one of its two incident cycle edges is traversed against the reference direction, and hence has $\sigma_C(e) = -1$, so that $t(e) < 0$. Thus every vertex on C benefits on at least one incident edge.

Now choose a forest spanning the vertices outside C, rooted at the vertices of C. For every vertex v outside C, let e_v denote the unique forest edge incident to v that lies on the path from v toward C (i.e., the edge joining v to its parent in the rooted forest), and set $t(e_v) > 0$ if $v \in V^A$ and $t(e_v) < 0$ if $v \in V^B$. Set $t(e) = 0$ on any remaining edges. Every vertex outside C then benefits on the edge joining it to its parent, while every vertex on C already benefits on an edge of C. Thus, t satisfies Conditions 1 and 2 of Lemma 1. ∎

Proof of Proposition 1. If: Let t have the described property. For each vertex $v^A \in V^A$, let $E_+(v^A) := \{e \in E(v^A) : t(e) > 0\}$ and $E_-(v^A) := \{e \in E(v^A) : t(e) \leq 0\}$. If $E_-(v^A) = \emptyset$, let μ^A_{ϵ,v^A} be uniform on $E(v^A)$. Otherwise, define

$$\mu^A_{\epsilon,v^A} := \begin{cases} \frac{1-\epsilon}{|E_+(v^A)|} & \text{for} \quad e \in E_+(v^A) \\ \frac{\epsilon}{|E_-(v^A)|} & \text{for} \quad e \in E_-(v^A) \\ 0 & \text{for} \quad e \notin E(v^A). \end{cases}$$

Trader A prefers t to no trade at v^A (i.e., conditional on $e \in E(v^A)$) given μ^A_{0,v^A}; by the continuity of expected payoffs in beliefs, A also prefers t given the beliefs μ^A_{ϵ,v^A} when ϵ is sufficiently small. Take μ^A to lie in the interior of the convex hull of $\{\mu^A_{\epsilon,v^A} : v^A \in V^A\}$. For ϵ sufficiently small, t provides strictly positive expected payoffs for A; additionally, by construction, μ^A is positive on every state, as required. To construct μ^B, repeat the argument for Trader B, assigning the higher probabilities this time to the states in each of B's information sets (i.e., states adjacent to B's vertices) for which $t(e) < 0$. The resulting beliefs μ^A and μ^B render t mutually agreeable.
Only If: Clearly if at some $v^A \in V^A$ we have $t(e) \leq 0$ for each $e \in E(v^A)$, then $\mathbb{E}^A[t(e)|e \in E(v^A)] \leq 0$, and Trader A does not strictly prefer t to no trade at v^A, which contradicts t being mutually agreeable. An analogous argument establishes Condition 2 for every $v^B \in V^B$. ∎

Proof of Lemma 2. Without loss of generality, assume that $W(\vec{C}) > 0$ and $\sigma_{\vec{C}}(e_1) = 1$. Let the states in the cycle be denoted by $C = \{e_1, e_2, \ldots, e_K\}$. Notice that $W(\vec{C}) > 0$ implies that $\sum_{e \in C} \sigma_{\vec{C}}(e)[\ln \mu^A(e) - \ln \mu^B(e)] > 0$, and hence

$$\Pi := \left(\frac{\mu^A(e_1)}{\mu^A(e_2)} \frac{\mu^A(e_3)}{\mu^A(e_4)} \cdots \frac{\mu^A(e_{K-1})}{\mu^A(e_K)} \right) \bigg/ \left(\frac{\mu^B(e_1)}{\mu^B(e_2)} \frac{\mu^B(e_3)}{\mu^B(e_4)} \cdots \frac{\mu^B(e_{K-1})}{\mu^B(e_K)} \right) > 1. \tag{C.1}$$

Choose $\delta \in (0, 1)$ such that $\delta^{K-1}\Pi > 1$.

The proof constructs a t that gives each Trader i positive expected value at each vertex $v \in V^i$. Throughout the remainder of the proof, it is convenient to work with unnormalized expected transfers. For a generic trade $\hat{t}$, define

$$M^i(v;\hat{t}) := \sum_{e\in E(v)} \mu^i(e)\hat{t}(e).$$

Since $\mu^i(E(v)) > 0$, $M^i(v;\hat{t})$ has the same sign as Trader i's conditional expected transfer at v. Thus Trader A strictly prefers $\hat{t}$ at v whenever $M^A(v;\hat{t}) > 0$, while Trader B strictly prefers $\hat{t}$ whenever $M^B(v;\hat{t}) < 0$.

We first construct an auxiliary trade, $\tilde{t}$, that is non-zero only at the states in C. Suppose $\tilde{t}(e) = 0$ for all $e \notin C$. For $e \in C$, we will define $\tilde{t}$ recursively as follows. Let $\tilde{t}(e_1) = 1$. Consider $s \in \{2, \ldots, K\}$. For s even, then

$$\tilde{t}(e_s) = -\delta\frac{\mu^A(e_{s-1})}{\mu^A(e_s)}\tilde{t}(e_{s-1});$$

for s odd, then

$$\tilde{t}(e_s) = -\delta\frac{\mu^B(e_{s-1})}{\mu^B(e_s)}\tilde{t}(e_{s-1}).$$

Given this construction, $\tilde{t}(e_s) < 0$ if s is even and $\tilde{t}(e_s) > 0$ if s is odd.

Note that for each even s, states e_{s-1} and e_s are both incident to the same vertex v_{s-1}, which is possessed by Trader A. Since $E(v_{s-1})$ contains no additional states in the cycle, we have

$$\begin{aligned} M^A(v_{s-1};\tilde{t}) &= \mu^A(e_{s-1})\tilde{t}(e_{s-1}) - \mu^A(e_s)\delta\frac{\mu^A(e_{s-1})}{\mu^A(e_s)}\tilde{t}(e_{s-1}) \\ &= (1-\delta)\mu^A(e_{s-1})\tilde{t}(e_{s-1}) > 0, \end{aligned} \tag{C.2}$$

which implies that A strictly prefers $\tilde{t}$ to no trade at each vertex (i.e., information set) they possess in the cycle.

Likewise, for each s that is odd with $2 < s < K$, states e_{s-1} and e_s are both incident to the same vertex v_{s-1}, which is possessed by Trader B, and $E(v_{s-1})$ contains no other states in the cycle. Thus,

$$\begin{aligned} M^B(v_{s-1};\tilde{t}) &= \mu^B(e_{s-1})\tilde{t}(e_{s-1}) - \mu^B(e_s)\delta\frac{\mu^B(e_{s-1})}{\mu^B(e_s)}\tilde{t}(e_{s-1}) \\ &= (1-\delta)\mu^B(e_{s-1})\tilde{t}(e_{s-1}) < 0. \end{aligned} \tag{C.3}$$

Notice that states e_K and e_1 are incident to the same vertex, v_K, which is possessed by Trader B, and iterating over our construction of $\tilde{t}$ on the cycle yields

$$\begin{aligned}
M^B(v_K;\tilde{t}) &= \mu^B(e_1)\tilde{t}(e_1) - \mu^B(e_K)\delta\frac{\mu^A(e_{K-1})}{\mu^A(e_K)}\tilde{t}(e_{K-1}) \\
&= \left(\mu^B(e_1) - \mu^B(e_K)\delta^{K-1}\frac{\mu^A(e_{K-1})}{\mu^A(e_K)}\frac{\mu^B(e_{K-2})}{\mu^B(e_{K-1})}\cdots\frac{\mu^A(e_1)}{\mu^A(e_2)}\right)\tilde{t}(e_1) \\
&= \left(\mu^B(e_1) - \delta^{K-1}\frac{\mu^B(e_K)}{\mu^A(e_K)}\frac{\mu^A(e_{K-1})}{\mu^B(e_{K-1})}\cdots\frac{\mu^B(e_2)}{\mu^A(e_2)}\mu^A(e_1)\right)\tilde{t}(e_1) \\
&= \left(1-\delta^{K-1}\Pi\right)\mu^B(e_1)\tilde{t}(e_1) \\
&< 0,
\end{aligned}$$

where the inequality follows from the fact that $\delta^{K-1}\Pi > 1$ and $\tilde{t}(e_1) = 1$. Thus, Trader B strictly prefers $\tilde{t}$ to no trade at each vertex (i.e., information set) they possess in the cycle.

The proof is complete when $C = E$. Suppose instead that $C \neq E$. We aim to construct a trade t that each Trader i prefers to no trade on *each* vertex they possess, including those that are not in the cycle. We construct such a t by augmenting $\tilde{t}$ such that t is approximately equal to $\tilde{t}$ on vertices in the cycle—hence preserving traders' strict preference for trade on those vertices as constructed above—yet also providing traders with an arbitrarily small, yet positive, expected payoff on vertices off the cycle.

Specifically, for each e_s in the cycle, let $t(e_s) = \tilde{t}(e_s)$. If each vertex of each trader is in the cycle, then take $t = \tilde{t}$, and we are done.

We now construct t for the alternative case where the cycle does not contain each vertex of each trader. First, consider a vertex v that (i) is part of the cycle, and (ii) each neighbor of v is also part of the cycle. For every such vertex v, take $t(e) = \tilde{t}(e)$ for all $e \in E(v)$. Next, we construct t for the remaining vertices to ensure each trader strictly prefers t to no trade using the following iterative procedure. If t is not entirely defined at the end of a step $L = 1, 2, \ldots$, move to the next step; otherwise, we are done.

Step 1: Consider Trader $i \in \{A, B\}$ and vertex $v^i \in V^i$ such that: (i) v^i is part of the cycle, and (ii) there exist $Z \geq 1$ neighbors of v^i that are not part of the cycle. For each such $v^j \in N(v^i)$, consider an arbitrary state in $E(v^i) \cap E(v^j)$ and label it e_{m_z}. This yields Z states $\{e_{m_1}, \ldots, e_{m_Z}\}$, each connecting v^i with a distinct vertex possessed by j (i.e., lying in a distinct information set of Trader j). For each $z = 1, \ldots, Z$, take $t(e_{m_z}) = \epsilon(v^i)$ where $\epsilon(v^i)$ is a constant such that $\epsilon(v^i) < 0$ if $i = A$ and $\epsilon(v^i) > 0$ if $i = B$; take $t(e) = 0$ for all other $e \in E(v^i) \setminus C$. Hence, $M^i(v^i;t) = M^i(v^i;\tilde{t}) + \epsilon(v^i)\sum_{z=1}^{Z}\mu^i(e_{m_z})$. Given inequalities (C.2) and (C.3), choose $\epsilon(v^i)$ small enough that $M^i(v^i;t)$ is positive if $i = A$ and negative if $i = B$. Thus, Trader i strictly prefers t to no trade at v^i, and Trader j earns a positive payoff at each e_{m_z}. Repeat this construction for all of Trader i's vertices that are part of the cycle yet have neighbors that are not part of the cycle. And then do the same for all such vertices of Trader $j \neq i$, swapping the roles of Trader i and j. If no such information sets exist for either player, we are done. Otherwise, move to Step 2 after Step 1.

Step $L \geq 2$: Let E_L denote all states at which t has so far been defined. If $E_L = E$, then we are done. Otherwise, consider Trader $i \in \{A, B\}$ and a vertex $v^i \in V^i$ such that: (i) $E(v^i) \cap E_L \neq \emptyset$, and (ii) there exist $Z \geq 1$ vertices possessed by Trader j such that $E(v^j) \cap E_L = \emptyset$. For each such neighboring vertex v^j, consider an arbitrary state in $E(v^i) \cap E(v^j)$ and label it e_{m_z}. This yields Z states $\{e_{m_1}, \ldots, e_{m_Z}\}$, each connecting v^i with a distinct vertex possessed by j (i.e., lying in a distinct information set of Trader j). For each $z = 1, \ldots, Z$, take $t(e_{m_z}) = \epsilon(v^i)$ where $\epsilon(v^i)$ is a constant such that $\epsilon(v^i) < 0$ if $i = A$ and $\epsilon(v^i) > 0$ if $i = B$; take $t(e) = 0$ for all other $e \in E(v^i) \setminus E_L$. Hence, $M^i(v^i; t) = M^i(v^i; t_{E_L}) + \epsilon(v^i) \sum_{z=1}^{Z} \mu^i(e_{m_z})$. By the previous step, $M^i(v^i; t_{E_L})$ is positive if $i = A$ and negative if $i = B$. Choose $\epsilon(v^i)$ small enough that $M^i(v^i; t)$ is positive if $i = A$ and negative if $i = B$. Thus, Trader i strictly prefers t to no trade at v^i. Repeat this construction for all of Trader i's vertices that have adjacent states in E_L yet have neighboring vertices that do not. And then do the same for all such vertices of Trader $j \neq i$, swapping the roles of Trader i and j. If no such vertices exist for either player, we are done. Otherwise, move to Step $L + 1$ after Step L.

This process will end in finite steps with t defined at all states. By construction, both traders prefer t to no trade at each of their vertices. ■

Proof of Proposition 2. Suppose, toward a contradiction, that the claim is false, and choose a counterexample G having the smallest possible number of edges. Let t be a mutually agreeable trade on G that does not follow the direction of disagreement along any oriented cycle $\vec{C}$ with $W(\vec{C}) \neq 0$.

By minimality, $t(e) \neq 0$ for every $e \in E$. To see this, suppose $t(e) = 0$ for some edge e. Delete e and let $E' := E \setminus \{e\}$. Restrict each trader's prior to E' and renormalize. Writing the renormalized priors as $\mu^{i\prime}(f) = \mu^i(f)/\mu^i(E')$ for $f \in E'$, the new edge weights satisfy $w'(f) = w(f) + \ln \mu^B(E') - \ln \mu^A(E')$. Thus renormalization adds a common constant to all remaining edge weights. Since $\sum_{f \in C} \sigma_{\vec{C}}(f) = 0$ for every oriented cycle $\vec{C}$ in the reduced graph, it follows that $W'(\vec{C}) = W(\vec{C})$. At either endpoint of e, the numerator of the trader's conditional expected payoff is unchanged, since $t(e) = 0$, while only its positive denominator changes. Hence the sign of the expected payoff is unchanged. Expected payoffs at all other vertices are unaffected. If deleting e leaves the graph connected, the restricted trade is therefore mutually agreeable on a graph with fewer edges. If deleting e disconnects the graph, the restricted trade is mutually agreeable on each resulting connected component. By Lemma 1, each such component therefore contains a cycle. Since every cycle in either component was already a cycle in G, either component gives a counterexample with fewer edges. In either case we contradict the minimality of G.

If $W(\vec{C}) = 0$ for every oriented cycle $\vec{C}$, then all of the cycle equations hold (see Equation 3), and the traders' interim beliefs are consistent with a common prior. This contradicts the existence of the mutually agreeable trade t. Hence there is an oriented cycle $\vec{C}$ with $W(\vec{C}) \neq 0$. Given this, let $\tilde{t}_C$ denote the trade supported on C constructed in the proof of Lemma 2. Thus $\tilde{t}_C$ strictly benefits the appropriate trader at every vertex incident to C and follows the direction of disagreement along $\vec{C}$. Since t does not follow the direction of disagreement along $\vec{C}$, and $t(e) \neq 0$ for every $e \in C$, there is some $e \in C$ for which $t(e)$ and $\tilde{t}_C(e)$ have opposite signs. Define

$$\lambda^* := \min_{\{e \in C : t(e)\tilde{t}_C(e) < 0\}} \frac{|t(e)|}{|\tilde{t}_C(e)|} > 0$$

and let $t^* := t + \lambda^* \tilde{t}_C$. Because $\tilde{t}_C$ strictly benefits the appropriate trader at every vertex incident to C and is zero elsewhere, t^* remains mutually agreeable. By construction, $t^*(e^*) = 0$ for at least one $e^* \in C$, while every nonzero component of t^* has the same sign as the corresponding component of t.

Delete one such edge e^* and let $E' := E \setminus \{e^*\}$. Since e^* belongs to a cycle, it is not a bridge, so deleting it leaves the resulting information graph connected. Restrict each trader's prior to E' and renormalize. As above, this renormalization shifts all remaining edge weights by a common constant and therefore leaves $W(\vec{C}')$ unchanged for every oriented cycle $\vec{C}'$ in the reduced graph. At either endpoint of e^*, the numerator of the trader's conditional expected payoff is unchanged because $t^*(e^*) = 0$, while its denominator remains positive. Thus the sign of the conditional expected payoff is unchanged. Expected payoffs at all other vertices are unaffected. Hence the restriction of t^* to the resulting information graph is mutually agreeable.

Finally, every oriented cycle of this smaller graph is also an oriented cycle of G. Consider any such $\vec{C}'$ with $W(\vec{C}') \neq 0$. If t^* is zero on some edge of C', then t^* does not follow the direction of disagreement along $\vec{C}'$. Otherwise, all of its signs on C' coincide with those of t, which by assumption did not follow the direction of disagreement along $\vec{C}'$. We have therefore constructed a counterexample with strictly fewer edges, contradicting the minimality of G. ■

Proof of Proposition 3. Part 1. By Lemma 2, if G admits no mutually agreeable trade, then for any oriented cycle $\vec{C}$, $W(\vec{C}) = 0$. Let $\widehat{G} = (\widehat{V}, E, \hat{\partial}_0, \hat{r}, w)$ be obtainable from $G = (V, E, \partial_0, r, w)$ by refining information. Let $h : \widehat{V} \to V$ send each refined vertex in $\widehat{G}$ to its parent vertex in G, namely $\hat{\partial}_0(e) = \{\hat{v}_1, \hat{v}_2\}$ implies that $\partial_0(e) = \{\, h(\hat{v}_1),\ h(\hat{v}_2)\}$. Refinement does not affect the edge disagreement $w(e)$.

Take any oriented cycle $\vec{\hat{C}} = (\hat{e}_1, \ldots, \hat{e}_K)$ in $\widehat{G}$. Mapping it into G by associating each vertex $\hat{v}$ with $h(\hat{v})$ while preserving its edges and orientation, $\sigma_{\vec{\hat{C}}}$, yields a closed, oriented walk $\vec{D}$ in G. By Veblen (1912)'s Theorem for directed graphs, $\vec{D}$ can be decomposed into a collection $\mathcal{C}$ of edge-disjoint oriented cycles, where each cycle inherits its orientation from $\vec{D}$. Therefore,

$$W(\vec{\hat{C}}) = \sum_{\vec{C} \in \mathcal{C}} \sum_{e \in C} \sigma_{\vec{C}}(e) w(e) = \sum_{\vec{C} \in \mathcal{C}} W(\vec{C}) = 0.$$

Applying Rodrigues-Neto (2009) Proposition 1 implies the existence of a common prior, which Milgrom and Stokey (1982) showed precludes mutually agreeable trade.

Part 2. For each $i = A, B$, let $V^i := r^{-1}(i)$ and $\widehat{V}^i := \hat{r}^{-1}(i)$. To show that t is mutually agreeable under G, we first show that $\mathbb{E}^A[t|e \in E(v)] > 0$ for each $v \in V^A$. Consider any $v \in V^A$. There exists a set of $J \geq 1$ vertices $\hat{v}_1, \hat{v}_2 \ldots, \hat{v}_J \in \widehat{V}^A$ such that $E(v) = E(\hat{v}_1) \cup E(\hat{v}_2) \cup \cdots \cup E(\hat{v}_J)$, where $E(\hat{v}_j) \cap E(\hat{v}_k) = \emptyset$ for $j \neq k$. Since t is mutually agreeable under each connected component of $\widehat{G}$, we have $\mathbb{E}^A[t|e \in E(\hat{v})] > 0$ for all $\hat{v} \in \widehat{V}^A$. This implies that $\mathbb{E}^A[t|e \in E(v)] > 0$ given that $E(v)$ is the union of the disjoint sets $E(\hat{v}_j)$. An analogous (omitted) argument implies that $\mathbb{E}^B[t|e \in E(v)] < 0$ for all $v \in V^B$. Thus, t is mutually agreeable under G. ■

Proof of Corollary 1. This follows from Part 2 of Proposition 3 when we take G to correspond to the case where neither trader has private information: for each Trader i, $V^i = \{v^i\}$ with $E(v^i) = E$. ■

Proof of Proposition 4. We prove Part 2 before Part 1.

Part 2. Let $t, t' \in T_{\text{BIN}}(S)$ for some $S \subsetneq E$. Since G and G' are connected, mutual agreeability implies ex-ante mutual agreeability. Moreover, any positive linear combination of two ex-ante mutually agreeable trades is itself ex-ante mutually agreeable. Define

$$t(e) = \begin{cases} x, & e \in S, \\ y, & e \notin S, \end{cases} \qquad t'(e) = \begin{cases} x', & e \in S, \\ y', & e \notin S. \end{cases}$$

Because each trade is mutually agreeable, its two values must have opposite signs. Toward a contradiction, suppose that t and t' do not have the same sign pattern. Without loss of generality, suppose $x > 0 > y$ and $x' < 0 < y'$.Choose $\lambda = -x/x' > 0$. Then

$$(t + \lambda t')(e) = \begin{cases} 0, & e \in S, \\ y + \lambda y', & e \notin S. \end{cases}$$

Since t and t' are both ex-ante mutually agreeable, $t + \lambda t'$ must also be ex-ante mutually agreeable. But a trade that is zero on S and constant on $E \setminus S$ cannot be strictly preferred to no trade by both traders: if $y + \lambda y' > 0$, Trader B dislikes it; if $y + \lambda y' < 0$, Trader A dislikes it; and if $y + \lambda y' = 0$, both are indifferent. This yields a contradiction. Hence t and t' must have the same sign pattern.

Part 1. Suppose, toward a contradiction, that $\operatorname{sgn}(t'(e)) = -\operatorname{sgn}(t(e))$ for every $e \in E$. Since t is binary, there exists some $S \subseteq E$ such that $t \in T_{\text{BIN}}(S)$. Mutual agreeability implies that the two values taken by t have opposite signs. Because t' is also binary and has the opposite sign from t in every state, its two values must be constant on the same two sets S and $E \setminus S$. Thus $t' \in T_{\text{BIN}}(S)$. Part 2 then implies that t and t' have the same sign pattern, a contradiction. ■

Proof of Proposition 5. Since G and G' are both connected and share the same edge set E, they represent two pairs of information partitions that share the same meet, $\{E\}$. Notice that if G is connected (i.e., it represents a pair of information partitions with a singleton meet equal to $\{E\}$), then a necessary condition for t to be mutually agreeable given information graph G is

$$\mu^A \cdot t > 0 \text{ and } \mu^B \cdot t < 0, \tag{C.4}$$

which follows from the fact that both players must strictly prefer t to no trade at the ex ante stage. If for some scalar $k \in \mathbb{R}$, $k\mathbf{1} - t$ is mutually agreeable under the given information graph G', then we must similarly have

$$\mu^A \cdot (k\mathbf{1} - t) > 0 \text{ and } \mu^B \cdot (k\mathbf{1} - t) < 0. \tag{C.5}$$

But since k is a scalar,

$$\mu^A \cdot t > 0 \text{ and } \mu^A \cdot (k\mathbf{1} - t) > 0 \Rightarrow k > 0, \tag{C.6}$$

and

$$\mu^B \cdot t < 0 \text{ and } \mu^B \cdot (k\mathbf{1} - t) < 0 \;\Rightarrow k < 0, \tag{C.7}$$

a contradiction. ■

Proof of Proposition 6. Since G is connected and contains a unique cycle, $|E| = |V|$. Because t is mutually agreeable, every vertex must benefit from t at some incident edge. Moreover, each nonzero edge can benefit exactly one of its two endpoints: if $t(e) > 0$, its A-endpoint benefits, whereas if $t(e) < 0$, its B-endpoint benefits. Since there are exactly as many edges as vertices, it follows that $t(e) \neq 0$ for every $e \in E$ and that every vertex benefits from exactly one incident edge.

We first characterize the signs of trade on the unique cycle C. By Proposition 2, t must follow the direction of disagreement along some oriented cycle $\vec{C}$ with $W(\vec{C}) \neq 0$. Since C is the unique cycle in G, this must be an orientation of C. By Definition 3, if $W(\vec{C}) > 0$, then $\sigma_{\vec{C}}(e)t(e) > 0$ for every $e \in C$, whereas if $W(\vec{C}) < 0$, then $\sigma_{\vec{C}}(e)t(e) < 0$ for every $e \in C$. Thus, following the direction of disagreement compactly implies $\operatorname{sgn}(t(e)) = \operatorname{sgn}(W(\vec{C}))\, \sigma_{\vec{C}}(e)$ for every $e \in C$.

Now consider the swapped information graph G^τ. Relative to the same traversal of C, swapping the traders' roles reverses the reference orientation of every edge. Hence the corresponding orientation satisfies $\sigma^\tau_{\vec{C}}(e) = -\sigma_{\vec{C}}(e)$ for every $e \in C$, and therefore $W^\tau(\vec{C}) = \sum_{e\in C} \sigma^\tau_{\vec{C}}(e)w(e) = -W(\vec{C})$. Applying Proposition 2 to the mutually agreeable trade t^τ under G^τ therefore gives

$$\text{sgn}(t^\tau(e)) = \text{sgn}(W^\tau(\vec{C}))\sigma^\tau_{\vec{C}}(e) = \text{sgn}(W(\vec{C}))\sigma_{\vec{C}}(e) = \text{sgn}(t(e))$$

for every $e \in C$. Thus, swapping information preserves the signs of trade on the cycle.

We next characterize the signs off the cycle. Each connected component of the subgraph formed by the edges outside C is a tree attached to C. Root each such tree at its unique vertex on C. As established above, every cycle vertex benefits from exactly one incident edge, and the sign pattern on C implies that this beneficial edge lies on the cycle. Hence every edge joining a cycle vertex to one of its children in a rooted tree must be unfavorable to the cycle vertex and therefore favorable to the child. The same reasoning applies inductively down each tree. Once a vertex outside C benefits from the edge joining it to its parent, that is its unique beneficial incident edge, so every edge joining it to one of its children must be unfavorable to it and favorable to the child. Thus, for every edge $e \in E \setminus C$, the endpoint farther from C benefits from $t(e)$.

Exactly the same argument applies to t^τ under G^τ: every off-cycle edge must benefit its endpoint farther from C. But swapping the role assignment switches whether that endpoint belongs to Trader A or Trader B. Since a positive trade benefits an A-vertex and a negative trade benefits a B-vertex, it follows that $\text{sgn}(t^\tau(e)) = -\text{sgn}(t(e))$ for every $e \in E \setminus C$.

We now turn to the mutual agreeability of the subtrade on C. At every vertex on C, all incident edges outside C are unfavorable to that vertex. Since t is mutually agreeable, removing those unfavorable off-cycle transfers leaves the trader at each cycle vertex strictly better off. Hence t_C gives strictly positive expected utility to the appropriate trader at every vertex incident to C. By Definition 5, it follows immediately that t_C is C-mutually agreeable under G. The same argument under the swapped role assignment shows that t^τ_C is C-mutually agreeable under G^τ.

Finally, suppose $E \setminus C \neq \emptyset$. At least one cycle vertex is incident to an edge outside C, and every such edge is unfavorable to that cycle vertex. Hence $t_{E\setminus C}$ gives the appropriate trader strictly negative expected utility at that vertex on the edge-induced subgraph $G_{E\setminus C}$. Thus $t_{E\setminus C}$ is not mutually agreeable on $G_{E\setminus C}$ and, by Definition 5, is not $(E \setminus C)$-mutually agreeable under G. The same argument applies to $t^\tau_{E\setminus C}$ under G^τ. ■

Proof of Proposition 7. Part 1. If: Suppose that t goes against the direction of disagreement along some oriented cycle $\vec{C}$. Let $\tilde{t}_C$ denote the cycle-supported trade constructed in the proof of Lemma 2 that follows the direction of disagreement along $\vec{C}$. Thus, $\tilde{t}_C$ strictly benefits the appropriate trader at every vertex incident to C and is zero outside C. Because t goes against the direction of disagreement along $\vec{C}$ while $\tilde{t}_C$ follows it, $t(e)\tilde{t}_C(e) < 0$ for every $e \in C$. Choose

$$\lambda \in \left(0, \min_{e\in C} \frac{|t(e)|}{|\tilde{t}_C(e)|}\right),$$

and define $\hat{t} := t + \lambda\tilde{t}_C$. For every $e \in C$, $\hat{t}(e)$ has the same sign as $t(e)$ and satisfies $|\hat{t}(e)| < |t(e)|$. Additionally, since $\hat{t}(e) = t(e)$ for every $e \notin C$, we have $\mathrm{sgn}(\hat{t}(e)) = \mathrm{sgn}(t(e))$ and $|\hat{t}(e)| \leq |t(e)|$ for every $e \in E$. Moreover, both traders weakly prefer $\hat{t}$ to t at every vertex and strictly prefer it at every vertex incident to C. Therefore t is interim inefficient.

Only If: Suppose that t is interim inefficient, and let $\hat{t}$ be a trade satisfying the conditions in Definition 6. Let $\Delta := \hat{t} - t$. Since $\hat{t}$ has the same sign as t and weakly smaller magnitude at every edge, for each $e \in E$, either $\Delta(e) = 0$ or $\Delta(e)$ has the opposite sign from $t(e)$. Moreover, because both traders weakly prefer $\hat{t}$ to t at every vertex, Δ is weakly preferred to no trade at every vertex, and it is strictly preferred to no trade by at least one trader at some vertex. Notice that the argument in the proof of Proposition 2 extends to a trade that is weakly preferred to no trade at every vertex and strictly preferred at some vertex.[36] Hence Δ follows the direction of disagreement along some oriented cycle $\vec{C}$ with $W(\vec{C}) \neq 0$, and $\Delta(e) \neq 0$ for every $e \in C$. Since $\Delta(e)$ has the opposite sign from $t(e)$ on each $e \in C$, t goes against the direction of disagreement along $\vec{C}$.

Part 2. Let $\bar{t}$ be any mutually agreeable trade, and define $\bar{u}^A(v^A) := \mathbb{E}^A[\bar{t}(e)|e \in E(v^A)]$ for each $v^A \in V^A$, and $\bar{u}^B(v^B) := \mathbb{E}^B[-\bar{t}(e)|e \in E(v^B)]$ for each $v^B \in V^B$. Mutual agreeability implies that all of these quantities are strictly positive. Consider the following maximization problem:

$$\max_t \left\{ \sum_{v^A \in V^A} \mathbb{E}^A[t(e)|e \in E(v^A)] + \sum_{v^B \in V^B} \mathbb{E}^B[-t(e)|e \in E(v^B)] \right\}$$

subject to

$$\begin{aligned} t(e)\bar{t}(e) \geq 0 \quad \text{and} \quad |t(e)| \leq |\bar{t}(e)| \qquad &\text{for every } e \in E; \\ \mathbb{E}^A[t(e)|e \in E(v^A)] \geq \bar{u}^A(v^A) \quad &\text{for every } v^A \in V^A; \\ \mathbb{E}^B[-t(e)|e \in E(v^B)] \geq \bar{u}^B(v^B) \quad &\text{for every } v^B \in V^B. \end{aligned}$$

The feasible set is nonempty because it contains $\bar{t}$. It is closed and bounded, and hence compact.

[36] To see this, restrict Δ to a connected component of the subgraph induced by the edges on which Δ is nonzero that contains a vertex at which one trader strictly benefits. At every vertex of this component, weak nonnegativity implies that the trader benefits on at least one incident edge. Lemma 1 therefore implies that this component contains a cycle. The minimal-counterexample argument in the proof of Proposition 2 can then be repeated with "mutually agreeable" replaced by "weakly beneficial at every vertex and strictly beneficial at some vertex." The same common-prior argument used in the proof of Proposition 2 to rule out cases in which all cycle equations hold also applies here, because the appropriate trader's expected payoff from Δ is weakly positive at every vertex and strictly positive at some vertex.

The objective function is continuous, so a maximizer t^* exists. Moreover, since $\bar{t}$ is mutually agreeable, t^* is mutually agreeable as well given that it weakly improves $\bar{u}^i(v^i)$ at each v^i for both traders $i = A, B$. We claim that t^* is interim efficient. Toward a contradiction, suppose otherwise, and let $\hat{t}$ render t^* interim inefficient according to Definition 6. Thus, for every $e \in E$, $\operatorname{sgn}(\hat{t}(e)) = \operatorname{sgn}(t^*(e))$ and $|\hat{t}(e)| \leq |t^*(e)|$, and both traders weakly prefer $\hat{t}$ to t^* at every vertex, with a strict preference at some vertex. These conditions imply that $\hat{t}(e)\bar{t}(e) \geq 0$ and $|\hat{t}(e)| \leq |\bar{t}(e)|$ for all $e \in E$ because t^* satisfies $t^*(e)\bar{t}(e) \geq 0$ and $|t^*(e)| \leq |\bar{t}(e)|$ for all $e \in E$. Moreover, because $\hat{t}$ weakly improves upon t^* at every vertex, it satisfies all of the expected-utility constraints in the maximization problem above. Thus $\hat{t}$ is feasible for that problem. Since at least one trader strictly prefers $\hat{t}$ to t^* at some vertex, $\hat{t}$ yields a strictly larger value of the problem's objective function, contradicting the optimality of t^*. Hence t^* is interim efficient. ∎

Proof of Proposition 8. By Proposition 2, t follows the direction of disagreement along some oriented cycle $\vec{C}$ in G. Since G and G' have common oriented cycles, $\vec{C}$ is also an oriented cycle in G' with the same orientation-sign pattern. Moreover, because the two graphs have the same edge weights w, the direction of disagreement along $\vec{C}$ is the same in both graphs. Since $\operatorname{sgn}(t'(e)) = -\operatorname{sgn}(t(e))$ for every $e \in E$, t' therefore goes against the direction of disagreement along $\vec{C}$ in G'. By Proposition 7, t' is interim inefficient.

Applying the same argument in the other direction, Proposition 2 implies that t' follows the direction of disagreement along some oriented cycle in G'. Since this oriented cycle is also common to G and G', the opposite-sign trade t goes against the direction of disagreement along that cycle in G. Hence, again by Proposition 7, t is interim inefficient. ∎

Proof of Lemma 3. The proof proceeds in two parts. First, we show the existence of a trade t that is interim efficient and mutually agreeable and whose signs induce an orientation that is strongly connected. Second, we show that on G^τ, the magnitude of t can be adjusted across states, without changing its sign in any state, to preserve interim efficiency and mutual agreeability.

Part 1. We denote the signs of the trade t by the orientation $\sigma_t : E \to \{\pm 1\}$, where $\sigma_t(e) = +1$ means that t's sign on e agrees with the reference orientation and A benefits; otherwise, B benefits. By Robbins (1939), because G is bridgeless and connected, there exists an orientation σ_t^0 that is strongly connected: for every $u, v \in V$, there exists a directed path under σ_t^0 from u to v.

Starting from σ_t^0, choose an oriented cycle $\vec{C}$ under σ_t^0 for which $W(\vec{C}) < 0$, if one exists. Reverse σ_t^0 on every edge in $\vec{C}$ to form σ_t^1. Doing this preserves strong connectivity.

Define

$$\Phi(\sigma_t) = \sum_{e \in E} \sigma_t(e) w(e).$$

Because $\vec{C}$ is an oriented cycle under σ_t, $\sigma_{\vec{C}}(e) = \sigma_t(e)$ for each $e \in C$. Therefore, whenever

$\sigma_t^1 \neq \sigma_t^0$,

$$\Phi(\sigma_t^1) - \Phi(\sigma_t^0) = \sum_{e \in C}(-\sigma_t^0(e)w(e) - \sigma_t^0(e)w(e)) = -2W(\vec{C}) > 0.$$

Repeat the process for each oriented cycle $\vec{C}$ under the current orientation for which $W(\vec{C}) < 0$. Since there are $2^{|E|}$ possible orientations, and Φ strictly increases at each step, the process must terminate after finite steps in an orientation that we call σ_t. This orientation is strongly connected and satisfies $W(\vec{C}) \geq 0$ for each directed cycle under σ_t.

Since σ_t is strongly connected, its oriented cycles span the cycle space of G (Proposition 6 of Gleiss, Leydold, and Stadler, 2003). Thus, if every oriented $\vec{C}$ under σ_t had $W(\vec{C}) = 0$, then by the linearity of W, every oriented $\vec{C}$ in G also has $W(\vec{C}) = 0$, which contradicts the assumption that $W(\vec{C}) \neq 0$ for some oriented cycle in G. Therefore, σ_t follows the direction of disagreement along some cycle $\vec{C}$ for which $W(\vec{C}) > 0$. Moreover, a trade whose signs agree with σ_t cannot go against the direction of disagreement on any oriented cycle. If a cycle is directed under σ_t, then its orientation has $W(\vec{C}) \geq 0$; if it is not directed under σ_t, the signs of the trade do not agree uniformly with either orientation of that cycle. Thus the trade goes against disagreement on no oriented cycle. By Proposition 7, any trade whose signs agree with σ_t is interim efficient. It remains to show that some such trades are mutually agreeable.

Let t_C follow the direction of disagreement, namely σ_t, along an oriented cycle $\vec{C}$ with $W(\vec{C}) > 0$; wlog, t_C gives positive utility to each vertex in G_C. (See the construction in the proof of Lemma 2.) We extend t_C to all of E by induction. Suppose that all $v \in V(S)$ benefits from trade t_S, where $V(S)$ are the vertices incident to edges in $S \subseteq E$; initially $V(S)$ is the vertices incident to edges in C. If $V(S) \neq V$, then since σ_t is strongly connected, there exists an edge e from some $u \in V(S)$ to some $v \in V \setminus V(S)$. Define a trade $t_{\{e\}}$ to have the sign $\sigma_t(e)$ on e, namely to give v positive utility. Since the trade $t_{\{e\}}$ harms u, scale up the existing trade t_S by some $\lambda > 0$ such that u strictly benefits from $t' = \lambda t_S + t_{\{e\}}$. Replace t with t' and $V(S)$ with $V(S) \cup \{v\}$. Repeating this procedure eventually yields a trade that each trader at each vertex prefers to no trade.

Finally, if t is zero on some edge e', assign $t(e')$ a sufficiently small value with the same sign as $\sigma_t(e')$ so as not to upend either trader's strict preference for trade at any vertex.

Part 2. Under the swapped role assignment, the reference orientation of every edge reverses. Thus, for any fixed traversal of a cycle, both $\sigma_{\vec{C}}$ and $W(\vec{C})$ reverse sign. Consequently, whether a fixed sign pattern follows or goes against the direction of disagreement is unchanged by the swap. Moreover, the orientation induced by the same trade signs under G^τ is the reverse of σ_t, and hence remains strongly connected. We can change the magnitudes of t_C to ensure the mutual agreeability at every vertex incident to an edge in C. We can then repeat the inductive argument of Part 1, changing the magnitude of trade but not its sign in any state, to ensure that every vertex in the graph benefits strictly from trade. The resulting trade is mutually agreeable and does not go against the direction of disagreement along any oriented cycle. By Proposition 7, Part 1, the resulting trade is therefore interim efficient under G^τ. ■

Proof of Proposition 9. Let S be a minimal-mutually-agreeable support, and let t_S be the trade associated with S in Definition 7. Notice first that $t_S(e) \neq 0$ for every $e \in S$. Otherwise, if $t_S(e) = 0$, then the restriction of t_S to $S \setminus \{e\}$ would remain mutually agreeable on $G_{S\setminus\{e\}}$, contradicting minimality.

We now establish two useful properties of G_S. First, G_S is connected. Suppose otherwise, and let R be the edge set of one of its connected components. Since t_S is mutually agreeable on G_S, its restriction t_R is mutually agreeable on G_R: the expected utility at every vertex of G_R is unchanged. Since $R \subsetneq S$, this contradicts the minimality of S.

Second, G_S is bridgeless. Suppose instead that $e' \in S$ is a bridge, with endpoints v^A and v^B. Removing e' separates G_S into two components. Let R_A and R_B denote the edge sets of the components containing v^A and v^B, respectively. Suppose first that $t_S(e') > 0$. Then e' is unfavorable to Trader B at v^B. Since t_S is mutually agreeable on G_S, v^B must be incident to another edge in S, so R_B is nonempty. Removing e' strictly increases Trader B's expected utility at v^B, while leaving expected utility unchanged at every other vertex of G_{R_B}. Hence t_{R_B} is mutually agreeable on G_{R_B}, contradicting the minimality of S. If $t_S(e') < 0$, the analogous argument applies to R_A. Hence G_S is bridgeless.

Next, because t_S is mutually agreeable on the connected graph G_S, Proposition 2 implies that t_S follows the direction of disagreement along some oriented cycle $\vec{C} \subseteq S$ for which $W(\vec{C}) \neq 0$.

We now decompose G along its bridges. Delete every bridge of G, and let $\mathcal{H}$ denote the resulting connected components, allowing a component to consist of a single isolated vertex. Contracting each component in $\mathcal{H}$ to a node and retaining the deleted bridges produces a tree, which we denote by T. Since G_S is connected and bridgeless, all edges in S lie in a single nontrivial component $H_0 \in \mathcal{H}$. Moreover, H_0 contains the cycle $\vec{C}$ identified above. Apply Lemma 3 to the subgraph induced by H_0. There exist trades t_{H_0} and $t^\tau_{H_0}$ that are mutually agreeable and interim efficient under the original and interchanged role assignments on H_0, respectively, and that have the same sign pattern on $E(H_0)$.

For each other nontrivial component $H \in \mathcal{H}$, choose a strongly connected orientation σ_H. Such an orientation exists because H is connected and bridgeless, as in the construction used in the proof of Lemma 3. We will use σ_H only to propagate surplus from the endpoint of the bridge connecting H to its parent component in T.

We now construct a trade on all of G. Root the tree T at H_0 and begin with the trade t_{H_0} on $E(H_0)$ that follows the direction of disagreement on $\vec{C}$ and is both mutually agreeable and interim efficient on G_{H_0}. Consider the bridges from H_0 to its child components. For each vertex v of H_0 incident to one or more such bridges, assign trade on those bridges so that each bridge benefits

its endpoint in the child component. Choose the magnitudes sufficiently small that, in total, these bridge trades reduce the expected utility of v by less than one-half of its current surplus. Thus every vertex of H_0 continues to receive strictly positive expected utility, while the endpoint in each child component receives strictly positive expected utility.

Proceed recursively down the tree T. Consider a component H that has received positive surplus at the endpoint v of its bridge from its parent. If H consists of a single vertex, no internal trade is required. Otherwise, use the strongly connected orientation σ_H to propagate this surplus throughout H. Begin with v as the set of reached vertices. Whenever not every vertex of H has been reached, strong connectivity implies that there is a directed edge under σ_H from some reached vertex u to an unreached vertex u'. Define trade on that edge with the sign prescribed by σ_H, so that u' benefits, and choose its magnitude sufficiently small that u continues to receive strictly positive expected utility. Add u' to the set of reached vertices and repeat. Since each step reaches a new vertex, this procedure terminates after finitely many steps with every vertex of H receiving strictly positive expected utility. Set trade equal to zero on all remaining edges of H.

The internal edges on which trade is nonzero form a spanning tree of H. Once every vertex of H receives strictly positive expected utility, assign trade on the bridges from H to its child components in the same way as above: each such bridge benefits its endpoint in the child component, and the magnitudes are chosen sufficiently small that the total loss at any vertex of H is less than one-half of its current surplus. Continuing recursively through the finite tree T defines a trade t on all of E that is mutually agreeable under G.

We next show that t is interim efficient. No bridge belongs to a cycle, so every cycle of G lies entirely within one component of $\mathcal{H}$. On H_0, the trade t agrees with the interim efficient trade t_{H_0} and therefore goes against the direction of disagreement on no oriented cycle. In every other nontrivial component H, the nonzero internal trades constructed above lie on a spanning tree of H. Hence every cycle in H contains at least one edge on which t is zero. Since going against the direction of disagreement along an oriented cycle requires the appropriate strict sign on every edge of that cycle, t cannot go against the direction of disagreement along any oriented cycle contained in H. Therefore t goes against the direction of disagreement on no oriented cycle in G. By Proposition 7, Part 1, t is interim efficient.

Finally, repeat the same construction under G^τ, beginning on H_0 with the trade $t^\tau_{H_0}$ supplied by Lemma 3. For each other nontrivial component H, choose a strongly connected orientation under the interchanged role assignment and propagate surplus from the endpoint of its parent bridge exactly as above, setting trade equal to zero on all internal edges not used in the propagation procedure. Continuing recursively through the same bridge tree yields a trade t^τ that is mutually agreeable under G^τ. As before, every cycle outside H_0 contains an internal edge on which t^τ is zero, while on H_0 the trade $t^\tau_{H_0}$ is interim efficient. Thus t^τ goes against the direction of disagreement on no oriented cycle in G^τ and is therefore interim efficient by Proposition 7, Part 1.

By Lemma 3, t_{H_0} and $t^\tau_{H_0}$ have the same sign pattern on $E(H_0)$. Since $S \subseteq E(H_0)$, the trades t and t^τ have the same sign pattern on S, as required. ■

Proof of Proposition 10. For each pair $(v^A, v^B) \in V^A \times V^B$, let $S(v^A, v^B) := E(v^A) \cap E(v^B)$. Construct a simple graph H on the vertex set V by placing a link between v^A and v^B whenever $S(v^A, v^B) \neq \emptyset$. Thus, H is obtained from G by replacing each nonempty collection of parallel edges between two vertices by a single link. Since G is connected, H is connected. Fix a spanning tree T of H. Consider any pair (v^A, v^B) whose corresponding link belongs to the spanning tree T, and write $S := S(v^A, v^B)$. By assumption, the weights are nonconstant on S, so there exist $e, e' \in S$ with $w(e) \neq w(e')$. The two edges $\{e, e'\}$ form a cycle in the two-vertex information subgraph G_S, and one of its orientations $\vec{C}$ satisfies $W(\vec{C}) = w(e) - w(e') \neq 0$. Hence, by Lemma 2, G_S admits a mutually agreeable trade. By Proposition 7, Part 2, there therefore exists a trade t^S on S that is both mutually agreeable and interim efficient on G_S.

Notice that t^S has the same properties when the roles of the two vertices are swapped. Indeed, in the two-vertex graph G_S, both vertices are incident to the same edge set S. Thus each trader's information set is S regardless of which vertex is assigned to that trader, so swapping the role assignment does not change either trader's conditional beliefs or preferences over trades on S. Define a trade t on E by

$$t(e) = \begin{cases} t^{S(v^A,v^B)}(e), & \text{if } e \in S(v^A, v^B) \text{ for some pair } (v^A, v^B) \text{ whose link belongs to } T, \\ 0, & \text{otherwise.} \end{cases}$$

The sets $S(v^A, v^B)$ are disjoint, so this defines t uniquely.

We first show that t is mutually agreeable under both G and G^τ. Fix any vertex $v \in V$. Since T is a spanning tree, v is incident to at least one link of T. For every pair (v^A, v^B) whose corresponding link belongs to T and is incident to v, the associated trade $t^{S(v^A,v^B)}$ is mutually agreeable on $G_{S(v^A,v^B)}$. In particular, $\sum_{e \in S(v^A,v^B)} \mu^A(e) t^{S(v^A,v^B)}(e) > 0$ and $\sum_{e \in S(v^A,v^B)} \mu^B(e) t^{S(v^A,v^B)}(e) < 0$. Summing these inequalities over all tree edges incident to v gives $\sum_{e \in E(v)} \mu^A(e) t(e) > 0$ and $\sum_{e \in E(v)} \mu^B(e) t(e) < 0$. Dividing by the corresponding positive probabilities of $E(v)$ does not change either sign. Thus Trader A strictly prefers t to no trade whenever v is assigned to A, and Trader B strictly prefers t to no trade whenever v is assigned to B. Since both inequalities hold at every vertex, t is mutually agreeable under both G and G^τ.

It remains to establish interim efficiency. By Proposition 7, Part 1, it suffices to show that t does not go against the direction of disagreement along any oriented cycle. First consider a cycle C that contains edges joining more than one pair of vertices. Collapsing each collection of parallel edges in C to a single edge produces a cycle in H. Since T is a tree, this cycle must contain some link of H that does not belong to T. Consequently, C contains some edge e belonging to a bundle $S(v^A, v^B)$ for a pair (v^A, v^B) whose corresponding link does not belong to T. By construction, $t(e) = 0$. Hence t cannot go against the direction of disagreement along $\vec{C}$: if $W(\vec{C}) > 0$, going against disagreement would require $\sigma_{\vec{C}}(e)t(e) < 0$ for every $e \in C$, while if $W(\vec{C}) < 0$, it would require $\sigma_{\vec{C}}(e)t(e) > 0$ for every $e \in C$. Since $t(e) = 0$ on at least one edge of C, neither condition can hold.

Next consider cycles consisting of two parallel edges. Such a cycle lies entirely within some $S(v^A, v^B)$. If the corresponding link between v^A and v^B does not belong to T, then t is identically zero on this bundle, so again it cannot go against disagreement on the cycle. If the corresponding link belongs to T, then t agrees on this bundle with the interim-efficient trade $t^{S(v^A,v^B)}$. By Proposition 7, Part 1, $t^{S(v^A,v^B)}$ does not go against the direction of disagreement along any cycle in G_S. Thus t does not go against disagreement along this cycle either.

We have therefore shown that t goes against the direction of disagreement on no oriented cycle in G. Proposition 7, Part 1, implies that t is interim efficient under G. The same argument applies under G^τ. For each pair (v^A, v^B) whose corresponding link belongs to T, the trade $t^{S(v^A,v^B)}$ remains interim efficient after the role assignment is swapped. Moreover, any cycle that contains edges from $S(v^A, v^B)$ for more than one pair (v^A, v^B) must contain an edge associated with a link of H that does not belong to T, and t is zero on every such edge. Thus t does not go against the direction of disagreement along any oriented cycle under G^τ. Thus, t is interim efficient under G^τ. ■

Proof of Proposition 11. Reverse the orientation of $\vec{C}$ if necessary so that $W(\vec{C}) > 0$. Fix any $\alpha \in (0, 1)$ and define a trade supported on C by

$$t(e) = \sigma_{\vec{C}}(e)[\mu^A(e)]^{-\alpha}[\mu^B(e)]^{-(1-\alpha)}$$

for $e \in C$ and $t(e) = 0$ for $e \notin C$.

We first show that, at every vertex incident to C, either trader would strictly prefer t to no trade if assigned to that vertex. Consider any such vertex. Its two incident cycle edges, which we denote by e_+ and e_-, satisfy $\sigma_{\vec{C}}(e_+) = 1$ and $\sigma_{\vec{C}}(e_-) = -1$. Since $W(\vec{C}) > 0$ and disagreement is monotone on $\vec{C}$, Definition 8 implies $w(e_+) - w(e_-) > 0$. Trader A's unnormalized expected transfer from the two cycle states at this vertex is therefore

$$\mu^A(e_+)t(e_+) + \mu^A(e_-)t(e_-) = \left(\frac{\mu^A(e_+)}{\mu^B(e_+)}\right)^{1-\alpha} - \left(\frac{\mu^A(e_-)}{\mu^B(e_-)}\right)^{1-\alpha} = e^{(1-\alpha)w(e_+)} - e^{(1-\alpha)w(e_-)} > 0.$$

Likewise, Trader B's unnormalized expected transfer is

$$\mu^B(e_+)t(e_+) + \mu^B(e_-)t(e_-) = \left(\frac{\mu^B(e_+)}{\mu^A(e_+)}\right)^\alpha - \left(\frac{\mu^B(e_-)}{\mu^A(e_-)}\right)^\alpha = e^{-\alpha w(e_+)} - e^{-\alpha w(e_-)} < 0.$$

Since Trader B's utility from the trade is the negative of the transfer, the latter inequality means that Trader B also strictly prefers t to no trade at this vertex. Thus, at every vertex incident to C, Trader A would strictly prefer t if that vertex were assigned to A, and Trader B would strictly prefer t if it were assigned to B. These conclusions therefore hold under both G and G^τ. It follows that t is mutually agreeable on the edge-induced subgraph G_C under either role assignment. Hence, by Definition 5, t is C-mutually agreeable under both G and G^τ.

Finally, we show that t is interim efficient under both information graphs. Since the factors $[\mu^A(e)]^{-\alpha}[\mu^B(e)]^{-(1-\alpha)}$ are strictly positive, $\text{sgn}(t(e)) = \sigma_{\vec{C}}(e)$ for every $e \in C$. Because $W(\vec{C}) > 0$, t therefore follows the direction of disagreement along $\vec{C}$. Consider any other cycle C' in G. If C' has a different edge set from C, then it contains some edge outside C. Since t is zero on every such edge, t cannot go against the direction of disagreement along any orientation of C', because going against disagreement requires the appropriate strict sign inequality to hold on every edge of the cycle. Hence t goes against the direction of disagreement on no oriented cycle in G. By Proposition 7, Part 1, t is interim efficient under G.

The same conclusion holds under G^τ. For a fixed direction of traversal around C, swapping the traders' roles reverses the reference orientation of every cycle edge, so both $\sigma_{\vec{C}}$ and $W(\vec{C})$ reverse sign. Consequently, the condition that t follows the direction of disagreement on C is unchanged. Every other cycle again contains an edge outside C on which t is zero. Thus t goes against the direction of disagreement on no oriented cycle under G^τ and is hence interim efficient under G^τ as well. ■

Proof of Corollary 2. Define the zero-sum trade $q := (t^A - t^B)/2$. Since the institutional trader's transfer is $t^C = -(t^A + t^B)$, we can write $q = t^A + t^C/2$ and $-q = t^B + t^C/2$. Because t is mutually agreeable under G, Trader A strictly prefers t^A to no trade at every A-vertex and Trader B strictly prefers t^B to no trade at every B-vertex. By assumption, both traders also assign strictly positive expected value to t^C at every vertex. Hence Trader A strictly prefers q to no trade at every A-vertex, while Trader B strictly prefers q to no trade at every B-vertex. Thus q is mutually agreeable under G.

Suppose, toward a contradiction, that $t' = (t^B, t^A) + (k1, -k1)$ is mutually agreeable under G^τ. Since $t'^A + t'^B = t^A + t^B$, the institutional trader's transfer is unchanged. Applying the same argument under G^τ, the zero-sum trade $q' := (t'^A - t'^B)/2 = -(t^A - t^B)/2 + k\mathbf{1} = -q + k\mathbf{1}$ is mutually agreeable under G^τ. But Proposition 5 implies that if q is mutually agreeable under G, then $-q + k\mathbf{1}$ cannot be mutually agreeable under G^τ for any $k \in \mathbb{R}$. This yields a contradiction. ■

Proof of Proposition A.1. First note that for any information set, the sign of a trader's conditional expected payoff is the same as the sign of the corresponding unnormalized expected payoff, since all priors have full support.

Part 1. Suppose that $E^A(e) = E^B(e) =: S$. If trade that favors Trader i at e is cursed rationalizable, then there exists a trade t with the appropriate sign at e such that both traders strictly prefer t to no trade conditional on S. Since $E^i(e) = S$ for all $e \in S$, the restricted trade t_S is mutually agreeable on G_S and favors Trader i at e. Conversely, suppose there exists a trade t_S that is mutually agreeable on G_S and favors Trader i at e. Extend t_S to all of E by setting trade equal to zero outside S. Since both traders' information set at e is S, both strictly prefer this trade to no trade at e. Hence it is cursed rationalizable at e and favors Trader i there.

Part 2. Suppose that $E^i(e) = \{e\}$. If a trade favors Trader $j \neq i$ at e, then Trader i receives a strictly negative payoff at e. Because Trader i knows that the state is e, their conditional expected payoff from the trade is strictly negative. Thus Trader i does not prefer the trade to no trade, so the trade cannot be cursed rationalizable at e.

Part 3. Without loss of generality, let $i = A$ and $j = B$, and write $I := E^A(e)$ and $J := E^B(e)$. We have $I \neq \{e\}$ and $I \neq J$. We construct a cursed rationalizable trade with $t(e) < 0$, so that the trade favors Trader B at e. Suppose first that $I \setminus J \neq \emptyset$, and choose $f \in I \setminus J$. Set $t(e) = -1$ and $t(f) = M$, and set $t(e') = 0$ at every other state. For M sufficiently large, $\sum_{e' \in I} \mu^A(e')t(e') = -\mu^A(e) + M\mu^A(f) > 0$, so Trader A strictly prefers the trade at e. Since $f \notin J$, $\sum_{e' \in J} \mu^B(e')t(e') = -\mu^B(e) < 0$, so Trader B also strictly prefers the trade at e. Hence the trade is cursed rationalizable at e and favors Trader B there.

It remains to consider the case $I \setminus J = \emptyset$. Since $I \neq J$, we then have $I \subsetneq J$. Because $I \neq \{e\}$, choose $f \in I \setminus \{e\}$, and because $I \subsetneq J$, choose $g \in J \setminus I$. Set $t(e) = -1$, $t(f) = M$, and $t(g) = -N$, and set trade equal to zero elsewhere. Choose M sufficiently large that $-\mu^A(e) + M\mu^A(f) > 0$. Since $g \notin I$, the choice of N does not affect Trader A's conditional expected payoff. We can therefore choose N sufficiently large that $-\mu^B(e) + M\mu^B(f) - N\mu^B(g) < 0$. Thus both traders strictly prefer the trade at e, while $t(e) < 0$. Hence trade that favors Trader B is cursed rationalizable at e. The case $i = B$ and $j = A$ is symmetric, completing the proof. ■