Queen's Economics Department Working Paper No. 852

Partial Provability in Communication Games

Barton L. Lipman Duane J. Seppi

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

4-1992

852

Partial Provability in Communication Games

by

Barton L. Lipman
Queen's University

and

Duane J. Seppi
Carnegia Mellon University

Current Draft

April 1992

Applications," *Bell Journal of Economics*, 12, 1981, pp. 380–391.

Milgrom, P., and J. Roberts, "Relying on the Information of Interested Parties," *Rand Journal of Economics*, 17, 1986, pp. 18–32.

Moore, J., and R. Repullo, "Subgame Perfect Implementation," *Econometrica*, 56, 1988, September 1191–1220.

Okuno–Fujiwara, M., A. Postlewaite, and K. Suzumura, "Strategic Information Revelation," *Review of Economic Studies*, 57, January 1990, pp. 25–47.

Palfrey, T., and S. Srivastava, "Nash Implementation Using Undominated Strategies," *Econometrica*, 59, March 1991, 479–501.

Spence, A. M., *Market Signaling*, Cambridge, MA: Harvard University Press, 1974.

Partial Provability in Communication Games

by

Barton L. Lipman
Queen's University

and

Duane J. Seppi
Carnegie Mellon University

Queen's University Discussion Paper #852

Current Draft
April 1992

We thank Ed Green, Debra Holt, Georg Noeldeke, Tom Palfrey, and seminar participants at the University of Toronto and at the 1990 Carnegie Mellon Conference on Political Economy for helpful suggestions. Of course, any errors are our responsibility. Support for this research was provided by a Sloan Foundation Grant to the Center for the Study of Public Policy at Carnegie Mellon University. This paper replaces "Signaling, Lobbying, and Political Debate."

ABSTRACT

This paper studies the ability of interested parties to communicate private information credibly to a decision maker in settings where their payoffs depend on the decision maker's action, but not on their own information *per se*. Examples of such situations arise in advertising, corporate control proxy contests, political debates, lobbying, and trials. Clearly, if an interested party can provide some unambiguous proof suppporting his claims, then it is possible for information to be revealed despite his incentive to lie. We show that even with only very limited provability and surprisingly weak assumptions on the interested parties' preferences, full revelation of information is possible in equilibrium. We also address the role of the number of interested parties, the relationship between their preferences, and simultaneous versus sequential games in facilitating communication.

REFERENCES

Austen–Smith, D., and J. Wright, "Competitive Lobbying for Legislators' Votes," University of Rochester working paper, 1990.

Banks, J., "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory*, 50, April 1990, 309–325.

Crawford, V., and J. Sobel, "Strategic Information Transmission," *Econometrica*, 50, November 1982, pp. 1431–1451.

Fishman, M., and K. Hagerty, "The Optimal Amount of Discretion to Allow in Disclosure," *Quarterly Journal of Economics*, May 1990, pp. 427–444.

Fudenberg, D., and J. Tirole, "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, April 1991, pp. 236–260.

Green, J., and J.-J. Laffont, "Partially Verifiable Information and Mechanism Design," *Review of Economic Studies*, 53, July 1986, pp. 447–456.

Grossman, S., "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 1981, pp. 461–483.

Harrington, J., "The Revelation of Information through the Electoral Process," Johns Hopkins University working paper, 1989.

Harris, M., and R. Townsend, "Resource Allocation Under Asymmetric Information," *Econometrica*, January 1981, pp. 33–64.

Kreps, D., and G. Ramey, "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55, November 1987, pp. 1331–1348.

Maskin, E., "Nash Equilibrium and Welfare Optimality," working paper, 1977.

Milgrom, P., "Good News and Bad News: Representation Theorems and

and consider the game induced by this inference rule. Clearly, the induced game is a zero–sum game — more specifically, we could assign payoffs in a fashion consistent with preferences in such a way that the game is zero–sum. Hence if there are any other equilibria, both senders must be indifferent between the equilibrium corresponding to σ and the alternative. Since every possible inference is degenerate, genericity implies that every equilibrium has s inferred.

So suppose that preferences are strictly conflicting and that there is no σ satisfying the condition of Proposition 8. We now show that there is no separating equilibrium. Suppose not. Let (σ, δ) be a separating equilibrium. Clearly, there must be some s and s' with $s \succ_1 s'$, $\sigma_1(s) \in M(s')$ and $\sigma_2(s') \in M(s)$. Since σ is one–to–one, either $\sigma_1(s) \neq \sigma_1(s')$ or $\sigma_2(s) \neq \sigma_2(s')$ or both. Without loss of generality, suppose that $\sigma_1(s) \neq \sigma_1(s')$.

Consider the deviation by sender 1 to $\sigma_1(s)$ in state s'. Since this is an equilibrium, we must have $s' \succeq_1 \delta(\sigma_1(s), \sigma_2(s'))$. By strictly conflicting preferences, then, $\delta(\sigma_1(s), \sigma_2(s')) \succeq_2 s' \succ_2 s$. Hence 2 deviates in state s to $\sigma_2(s')$, a contradiction. ▌

I. Introduction.

A wide range of economic, political and legal processes involve efforts by one or more interested parties to persuade an uncommitted decision maker. Managers issue audited earnings reports to shareholders. Lobbyists provide legislators with data on different proposals. Litigants hire experts to testify on the facts of a case. One common feature of these examples is that frequently what each interested party (or speaker) wants the decision maker to do may be independent of his actual private information. The difficulty this creates is that it is then unclear what relation his statements bear to his private information and hence what information content they have.

Certainly, if a statement contains irrefutable proof of some fact, then it has information content regardless of the speaker's preferences. Unfortunately, how much an interested party can explicitly prove is often limited. That is, he may be able to prove some — but not necessarily all — of what he knows. We call this ***partial provability.*** Limitations on the ability to prove claims arise from at least two sources. First, it may simply not be possible to prove a fact definitively even if it is true either because proof does not exist or because speakers are not allowed to provide it. For example, while a pianist can easily demonstrate that he can play the piano, it is hard to imagine how a non-musician could prove that he cannot.[1] Alternatively, the exclusionary rule in trials prohibits the introduction of evidence (factual "proofs") which prosecutors do have but which was obtained through illegal searches. Second, there may be limitations on the number of facts speakers can disclose. For example, time limits may constrain information transmission if the decision maker is unable to absorb more than a certain amount of information in the time available.

---

[1] We thank Michael Peters for this example. In terms of the analysis that follows, this limitation is a consequence of the fact that everything a non-musician can do at the keyboard is a subset of what a pianist can do.

This paper investigates the possibilities for informative communication by interested parties in settings with partial provability where speakers' preferences are independent of their private information. In particular, we show how the decision maker's inference rule can exploit conflicts and commonalities in different speaker's preferences over his possible actions as a way of overcoming obstacles to communication due to limits on provability. The analysis is encouraging in that we find that a surprising amount of private information can often be revealed even when very little can be proven. The following example illustrates this point.

Example 1.

Suppose there are a large number — say 1000 — of possible "states of nature" and that in each state the decision maker (if he knew the true state) would take a different action. Suppose further that initially the decision maker is uninformed and must rely for guidence on two lobbyists who each know the true state, but who have conflicting preferences (*i.e.,* they have the opposite rankings over any pair of possible by the decision maker). Each lobbyist has one opportunity to speak, first lobbyist 1, then lobbyist 2. The available messages allow the lobbyists (1) to assert unverifiably that a particular state s is true (even if it is not) and then (2) to submit a single piece of evidence ruling out any single untrue state s' (*i.e.,* a state can not be ruled out if it is actually true). We refer to such messages as "not" messages since they unambiguously prove only that some single s' is not the true state. Thus, a "not" message is the minimally informative message in terms of what it explicitly proves.

Taken together the two lobbyists' messages definitively rule out only two of the thousand states. However, despite the limited informativeness of these message, there is an inference rule which supports a perfectly revealing equilibrium! The inference rule is simply to believe lobbyist 1's assertion unless it is disproven by lobbyist 2. In this case, believe the state (out of those not ruled out by the two messages) which leads to the worst possible

consistent with feasibility. If m_1 is an equilibrium message for 1 but m_2 is not an equilibrium message for 2, then $\delta(m_1, m_2)$ is the worst degenerate inference for sender 2 consistent with feasibility. Finally, suppose m_1 and m_2 are equilibrium messages but not for the same state. If there is any s' such that $m_1 = \sigma_1(s')$ and $m_2 \in M(s')$, then $\delta(m_1, m_2)$ is the best such s' for sender 1. Otherwise, it is the worst degenerate inference for sender 1 consistent with feasibility.

It is not hard to see that (σ, δ) is a separating equilibrium. Consider any deviation by sender 1 to message m_1 in state s. First, suppose that $(m_1, \sigma_2(s)) = \sigma(s')$ for some $s' \neq s$. Clearly, we must have $\sigma_2(s) = \sigma_2(s')$, so that $\sigma_2(s) \in M(s')$ and $\sigma_2(s') \in M(s)$. Hence by the condition of the theorem, it must be true that $\sigma_1(s') \in M(s)$ only if $s \succ_1 s'$. Hence the deviation makes sender 1 strictly worse off. Hence we may as well suppose that $(m_1, \sigma_2(s)) \neq \sigma(s')$ for any s'.

If m_1 is not an equilibrium message for 1, then, by construction, δ is the worst degenerate inference for 1 consistent with feasibility, so that $\delta \preceq_1 s$. By genericity, either $\delta = s$ or $\delta \prec_1 s$. Hence either the deviation has no effect or makes sender 1 strictly worse off.

Finally, suppose that sender 1 deviates to $m_1 = \sigma_1(s')$ for some s' but that $(\sigma_1(s'), \sigma_2(s)) \neq \sigma(s'')$ for any s''. Clearly, this requires $\sigma_1(s') \in M(s)$. Hence either $s \succ_1 s'$ or $\sigma_2(s) \notin M(s')$. Let $\hat{S}$ denote the set of s' such that $m_1 = \sigma_1(s')$ and $\sigma_2(s) \in M(s')$. If $\hat{S}$ is nonempty, then $s \succ_1 s'$ for every $s' \in \hat{S}$. Since the receiver's inference is such an s', sender 1 again is strictly worse off for having deviated. Alternatively, if $\hat{S}$ is empty, by similar arguments to the above, either the deviation has no effect or makes sender 1 strictly worse off. A similar argument shows that every possible deviation by sender 2 either has no effect or makes sender 2 worse off. Hence we have a separating equilibrium.

To see that this is a strong separating equilibrium, fix any state s

librium. Let $\delta^*(s)$ denote the receiver's inference in state s in equilibrium. We show by induction that $\delta^*(s) = s$ for all s. Without loss of generality, number the states so that $s_1 \succ_1 s_2 \succ_1 \ldots \succ_1 s_\ell$. (By genericity, all these comparisons are strict.) By (4), there is a message, say m_1, which is in $M(s_1)$ but not $M(s_i)$ for any $i \neq 1$. Since the receiver's inference rule must be consistent with feasibility, we must have $\delta(m_1) = s_1$. By betweenness, if $\delta^*(s_1) \neq s_1$, we must have $s_1 \succ_1 \delta^*(s_1)$. But then sender 1 could deviate in state s_1 to m_1, a contradiction. Hence we must have $\delta^*(s_1) = s_1$.

Suppose that for $k = 1, \ldots, K-1$, we have shown that $\delta^*(s_k) = s_k$. We complete the proof by showing that $\delta^*(s_K) = s_K$. Suppose not. Clearly, $\delta^*(s_K)$ cannot put any probability on states $s_1, \ldots, s_{K-1}$. Then betweenness implies that $s_K \succ_1 \delta^*(s_K)$. Furthermore, for any δ putting probability 1 on $\{s_1, \ldots, s_K\}$, betweenness implies $\delta \succeq_1 s_K \succ_1 \delta^*(s_K)$. By (4), there exists a message $m_K \in M(s_K)$ which is not in $M(s_i)$ for any $i \geq K+1$. Since the receiver's inference rule must be consistent with feasibility, the support of $\delta(m_K)$ must be contained in $\{s_1, \ldots, s_K\}$. Hence sender 1 would be better off deviating to m_K in state s_K, a contradiction. ∎

Proof of Proposition 8.

First, suppose that σ satisfies the condition of Proposition 8. To see that σ must be invertible, suppose, to the contrary, that there are two states, s and s', such that $\sigma(s) = \sigma(s')$. By genericity, we may assume that $s \succ_1 s'$. Hence by the condition of Proposition 8, either $\sigma_1(s) \notin M(s')$ or $\sigma_2(s') \notin M(s)$. Either way, we cannot have $\sigma_1(s) = \sigma_1(s')$ and $\sigma_2(s) = \sigma_2(s')$. Hence σ is one–to–one.

We now show that there is a δ such that (σ, δ) is a strong separating equilibrium. To construct this inference rule, say that m_i is an *equilibrium message* for i if there exists an s such that $\sigma_i(s) = m_i$. Clearly, for any s, $\delta(\sigma_1(s), \sigma_2(s)) = s$. For any (m_1, m_2) such that m_1 is not an equilibrium message for 1, let $\delta(m_1, m_2)$ be the worst degenerate inference for sender 1

action from 1's point of view. In equilibrium with this rule, lobbyist 1 tells the truth in every state (except possibly in his least preferred state where lying and getting caught does not affect the final outcome anyway). If 1 does not tell the truth, the availability of a complete set of "not" messages insures that lobbyist 2 can refute any false assertion by 1. This inference rule along with the conflicting preference assumption insures that 2 will want to do so.

Intuitively, complete revelation is possible because 2's message conveys more information in equilibrium than simply ruling out a single state. In particular, the failure of 2 to "refute" 1's claim is taken as evidence of the inability to refute. More generally, an interested party's failure to provide proofs of certain facts may be construed to mean that these facts are untrue. While this can occur with only one speaker, this example suggests that with multiple speakers, it can greatly increase the amount of information revealed. This is true because with multiple speakers, the interpretation of a message from one speaker can depend on what the other speakers say. Thus in the example, the interpretation of a "not state s'" message from 2 changes dramatically depending on whether 1 initially claimed state s' or not.

While there is a large literature on signaling and mechanism design, relatively little directly addresses the issues we study. Milgrom [1981], Milgrom and Roberts [1986] and Grossman [1981] study signaling in the special case of *complete provability* — that is, when an interested party can prove any (true) claim. As a result, an interested party can — and in equilibrium does — reveal all of his information. This approach has two weaknesses. First, it makes strong assumptions about the set of available messages. In particular, to prove the truth unambiguously may require exhaustive specificity in some of the proofs. Second, this approach cannot explain the prevalence of adversarial debate in many real world decision making processes. In particular, with complete provability and symmetrically informed parties, all information can be elicited from a single interested party. Thus there is no gain from informational competition among multiple interested parties.

Another polar case is to assume that no proof is possible. There are two branches of this literature. In one, all statements are completely unverifiable "cheap talk." This approach is taken in virtually all of the work on mechanism design — see, for example, Harris and Townsend [1981], Maskin [1977], Moore and Repullo [1988], and Palfrey and Srivastava [1991]. However, unlike the situations of interest here, these models assume that each interested party's preferences depend on his private information. Hence, in designing optimal institutions, one can exploit differences in preferences across different states of the world to generate the desired outcome as a function of the state. However, as Crawford and Sobel [1982] show, at least in the case where there is only one interested party, if his preferences are independent of his private information, "communicative" equilibria do not exist. The other branch, typified by Spence [1974], assumes that signals inherently prove nothing but have different costs depending on the signaller's private information. This induces different preferences for the signaller in different states of the world so that, just as above, information can be credibly communicated.[2]

The idea of partial provability[3] first appeared in Milgrom [1981]. The

---

[2] Recent papers by Banks [1990], Harrington [1989] and Austen-Smith and Wright [1990] investigate communication without provability in the context of political campaigns and lobbying. Again, some form of dependence between the informed's preferences and information is used to generate communicative equilibria (e.g., via an exogenous cost to lying or by assuming that the post-election cost of implementing a policy depends on the policy's "distance" from the electorate's preferences).

[3] The term "verifiability" is generally used in place of "provability." "Verifiability" suggests that the onus is on the listener to confirm the truth of the statement, while "provability" is intended here to suggest that it is on the speaker.

some s, the receiver infers s. For any other message, the receiver's inference is the worst inference for sender 1 consistent with feasibility.

A separating equilibrium exists with this inference rule. To see this, suppose, to the contrary, that in some state s, sender 1 has an optimal message, say m', leading to a strictly better inference than s. Suppose that $m' = \sigma(s')$ for some $s' \neq s$, so that $\delta(m') = s'$. If $m' \in M(s)$, then, by construction, $s \notin L_1(s')$ — that is, $s \succ_1 s'$, a contradiction. Hence it must be true that $m' \neq \sigma(s')$ for any s'. Therefore, $\delta(m')$ is the worst inference for sender 1 consistent with feasibility. Clearly, then $\delta(m') \preceq_1 s$, again a contradiction.

Now suppose preferences are generic. We show that a strong separating equilibrium exists iff (4) holds. First, suppose (4) does hold and consider the inference rule δ constructed above. We claim that every equilibrium in the induced game has the receiver inferring the truth. The arguments above establish that if in state s, sender 1 uses a message $m' = \sigma(s')$ for $s' \neq s$, then the receiver's inference is strictly worse for sender 1 than s. Any other deviation leads to the worst degenerate inference for him consistent with feasibility. Clearly, this inference, say s'', must satisfy $s'' \preceq_1 s$. If $s'' \neq s$, then, by genericity, $s'' \prec_1 s$, completing the argument.

Now suppose that (4) does not hold. We claim that no separating equilibrium, strong or otherwise, exists. Suppose not. Let σ be sender 1's strategy in a separating equilibrium. Fix any s such that (4) fails. Since (4) fails, there is some state s' such that $\sigma(s) \in M(s')$ and $s' \in L_1(s)$ — that is, $s' \preceq_1 s$. By genericity, then, $s' \prec_1 s$. Consider, then, the deviation in state s' to message $\sigma(s)$. This message is feasible and leads to a strictly preferred inference, contradicting the optimality of σ. ∎

Proof of Proposition 7.

Suppose the conditions of the Proposition hold and consider any equi-

The inference rule is as follows. Let sender 1's claim be s^1 and the vector of messages be $\underline{m}$. If $s^1 \in F(\underline{m})$, then the inference is s^1. Otherwise, the inference is the worst degenerate inference for sender 1 consistent with feasibility.

We claim that for any s, there is no equilibrium in the induced game given this inference rule such that the receiver's inference differs from s. Suppose, to the contrary, that such an equilibrium exists. First, suppose that the inference in this equilibrium is the same as sender 1's claim $s^1 \neq s$. Clearly, if sender 1 had claimed s, this would be the inference since it could not be refuted. Therefore, if this is an equilibrium, it must be true that $s^1 \succeq_1 s$. By genericity, we can assume this preference is strict. By construction of sender 1's preferences, then, it cannot be true that $M(s) \subset M(s^1)$. Also, since senders 1 and 2 have conflicting preferences, we have $s \succ_2 s^1$. Furthermore, if sender 2 used a message $m \in M(s) \setminus M(s^1)$, the resulting inference, say δ, would necessarily satisfy $\delta \preceq_1 s$ and hence $\delta \succeq_2 s \succ_2 s^1$. Clearly, such a message would be better for sender 2 than his "equilibrium" message, a contradiction.

Hence if s is not inferred in equilibrium, it must be true that sender 1's claim is refuted in equilibrium. But then the receiver's inference, say δ, satisfies $s \succeq_1 \delta$. Recall, though, that δ is a degenerate inference. Since it is not equal to s by assumption, genericity implies $s \succ_1 \delta$. But sender 1 could have claimed s, so the claim s^1 is not optimal, a contradiction. ▌

Proof of Proposition 6.

We first show that a separating equilibrium exists if (4) holds. Suppose (4) holds and construct any invertible function $\sigma : S \to M$ such that

$$\sigma(s) \in M(s) \setminus \bigcup_{s' \in L_1(s)} M(s'), \quad \forall s.$$

By (4) and our cardinality assumption, such a function must exist. Construct the receiver's inference rule as follows. For any message $m = \sigma(s)$ for

attraction of this approach is that while interested parties may not be able to prove much, it is often unrealistic to assume that they can prove *nothing*. The key insight is that limitations on provability can be overcome by the fact that far more may be inferred in equilibrium than is explicitly proved. In particular, as in Example 1, the failure to prove certain supporting facts may be construed as evidence that certain possible claims are untrue.

The subsequent literature on partial provability is still relatively small. Fishman and Hagerty [1990], like Milgrom [1981], consider the case of a single interested party in the context of a particular message structure. We consider multiple interested parties with a general message structure. Okuno–Fujiwara, Postlewaite, and Suzumura [1990] also consider multiple interested parties, again with a specialized message structure. Unlike our model, they assume that the interested parties are asymmetrically informed. Finally, Green and Laffont [1986] consider the role of partial provability in a principal–agent context.

This paper is organized as follows. We present the basic model in Section II. Rather than focus on a specific institutional setting, we consider a very general message environment. In Section III, we consider a simple sequential game, called an *open forum*, in which each interested party has one turn to speak. In particular, we focus on when separating equilibria — equilibria where the decision maker learns the state exactly — are possible. Furthermore, we focus on separating equilibria which are robust in the sense that fixing the inference rule, separation is the unique equilibrium outcome in the induced game played by the speakers. We call such equilibria *strong separating equilibria*. Hence our results are weaker than the uniqueness results in Milgrom and Roberts [1986] but stronger than simple existence.

We have the following results. First, weak preference assumptions often suffice to sustain a strong separating equilibrium when complete provability is relaxed to partial provability. In particular, as Example 1 shows, very limited provability is often enough. Second, given any message technology in which different states have different (but otherwise arbitrarily overlapping) sets of feasible messages, it is always possible to specify preferences for two (or more) interested parties which support a strong separating equilibrium. Two, however, is the minimum number of speakers for which this is true. With only one speaker, separation may not be possible for any preference specification. Third, we show that the inference rules which support strong separating equilibria — although more complicated than those used in complete provability models — often have intuitive interpretations. Lastly, we show by example that information revelation can be decreasing in the number of speakers, that uninformative messages may be a crucial part of a separating equilibrium, and that common interests among the speakers can encourage information revelation more than certain conflicting interests.

In Section IV we briefly contrast open forums with an alternate game in which interested parties send messages simultaneously to the decision maker. We show that neither an open forum nor the simultaneous game dominates the other in terms of their capacity to elicit information.

the induced game given this inference rule such that the receiver's inference is incorrect. Suppose, contrary to the claim, that there is an equilibrium in state s in which the receiver's inference is $\delta \neq s$. By construction, the receiver's inference is necessarily degenerate. Hence $\delta = s'$ for some $s' \neq s$. By conflicting preferences, there is a sender i for whom $s \succ_i s'$. At sender i's turn, there is some "claim on the table," say s''. Clearly, it must be true that $s'' \neq s$ — since s is true, it cannot be refuted, so if it is ever the claim on the table, it will be inferred by the receiver. Furthermore, since s'' had to be claimed in a trustworthy fashion, it cannot be true that $M(s) \subset M(s'')$. By (M4), there exists some $m \in M(s)$ such that m refutes s'' and every $\hat{s}$ such that $M(\hat{s}) \subset M(s)$. Hence sender i can use this message and claim s. This claim cannot be refuted, so s will be inferred. Since $s \succ_i s' = \delta$, this deviation is profitable for sender i, contradicting the hypothesis that we had an equilibrium. ▌

Proof of Proposition 5.

Choose any generic $\succ_1$ such that $M(s) \subset M(s')$ implies $s \succ_1 s'$ and let $\succ_2$ satisfy $\delta \succ_2 \delta'$ iff $\delta' \succ_1 \delta$. To see that this is possible, let P^1 denote the set of pairs (s, s') such that $M(s) \subset M(s')$. Clearly, the relation P^1 is asymmetric and transitive. If for every s, s', either $(s, s') \in P^1$ or $(s', s) \in P^1$, we are done. If not, choose any s, s' such that neither ordered pair is in P^1 and let P^2 be the transitive closure of $P^1 \cup \{(s, s')\}$. It is not hard to show that P^2 will also be asymmetric and transitive. Again, if it satisfies genericity, we are done; otherwise, we can continue this process. Since S is finite, this process necessarily generates an asymmetric and transitive relation, say P, such that sPs' or $s'Ps$ for every $s \neq s'$. Extend this relation to nondegenerate inferences in any manner consistent with weak preference being complete, reflexive, and transitive[15] and call the resulting (strict) preference $\succ_1$.

[15] This is obviously possible. For example, one can choose any of the degenerate inferences and assume that all nondegenerate inferences are indifferent to this inference.

generic, a separating equilibrium exists. ▌

Proof of Proposition 4.

For simplicity, we describe the inference rule as though senders send both a message $m \in M$ and claim a state $s \in S$. As discussed in the text, the cardinality condition implies that this is without loss of generality. Also, notice that the receiver is free to choose how to interpret a given claim s. That is, he can treat a claim of s as a claim of a different state s'. In this sense, it is without loss of generality to suppose that the receiver can force the senders to only make claims with certain properties as long as given any message, there is at least one claim satisfying these properties.

With this in mind, consider the following restrictions on claims. Let (m_i, s^i) denote sender i's message and claim, where we require throughout that sender i's claim must be a state not ruled out by the messages of the first i senders. We also require the claim s^i to have the property that m_i rules out every state s' not ruled out previously such that $M(s') \subset M(s^i)$ — that is, the claim must be trustworthy. We will refer to a claim satisfying these two criteria as a *valid* claim. To show that this restriction is legitimate, we show that given any possible messages by the first i senders, there is always at least one valid claim s^i. To see this, fix any i messages which are all feasible in at least one state. Let $F = \cap_{j=1}^{i} F(m_j)$. Clearly, F is finite and nonempty. Hence the fact that $\subset$ is transitive and asymmetric implies that there is at least one state $s' \in F$ such that there is no $s'' \in F$ with $M(s'') \subset M(s')$. This state would be a valid claim.

The receiver's inference rule is as follows. If $s^1 \in F(\underline{m})$, then the receiver infers s^1. If not, then let k denote the first i such that $s^1 \notin F(m_i)$. If $s^k \in F(\underline{m})$, then the receiver infers s^k. Etc. It is easy to see that this algorithm generates a feasible inference given any $\underline{m}$.

To conclude, we show that for any state s, there is no equilibrium in

II. The Model.

We have $n+1$ players, n "senders" of information and one "receiver." The senders send messages to the receiver, who then chooses an action that affects his utility and the utility of each sender. Each sender has information that the receiver does not have which affects the receiver's payoff, but not his own. Thus the sending of messages can be interpreted as the senders' attempts to persuade the receiver to choose an action they would like. We assume that the senders are symmetrically informed.

More formally, S is the set of possible states of the world. For simplicity, S is assumed to be finite. A given state $s \in S$ specifies all facts known to the senders which the receiver would like to find out.[4] For example, a state might specify the circumstances of a crime, the relative merits of various brands of a good, the talent of a firm's incumbent management, or the costs and benefits of various possible defense projects.

Senders inform the receiver by sending messages which may include evidence such as certain statements, documents, or physical items. Thus a state s also specifies the availability of such evidence. It is precisely the fact that these items are available in certain situations and not in others that makes them useful as evidence. For example, if the underlying facts are that the murder victim was stabbed and not shot, then there cannot be a bullet in the corpse. Less dramatically, a (legitimate) house deed with a given individual's name on it exists if and only if that person is a homeowner. Hence these items can act as evidence in the sense that showing them to the receiver proves to him that certain states are not possible.[5]

---

[4] In other words, we do not distinguish between states which senders cannot distinguish.

[5] This formulation also allows the possibility of messages that give noisy information. For example, if it is possible to forge a house deed, then if shown a house deed with a specific individual's name on it, the receiver rules out all states except those in which

There are at least two distinct forms of provability that our modeling approach can represent. First, as noted above, a message may include a presentation of documents or other "hard" evidence substantiating some set of facts. Second, a message can include a logical proof that known facts imply a particular conclusion. Such a proof cannot be produced if the facts are inconsistent with this claim.[6] Either way, using a message m feasible only in states in a set $\hat{S}$ proves that the true state must be in $\hat{S}$. Equivalently, such a message rules out every $s \notin \hat{S}$. We refer to this as the *pure information content* of m. Letting M denote the set of all possible messages, we have a function $M(s)$ which for each s gives that (nonempty) subset of M corresponding to those messages which are feasible in state s. It is also useful to define the inverse function, F, given by $F(m) = \{s \mid m \in M(s)\}$, the set of states in which the message m is feasible.

As Example 1 illustrates, messages can have information beyond what they prove. We refer to this as *equilibrium information content.* Pure information content derives directly from provability: that is, a message m at least proves $s \in F(m)$. Equilibrium content derives from the fact that the equilibrium strategies of the senders call for certain messages to be sent in certain states. Hence if sender i's strategy is to use message m only in those states in a subset $\hat{S}$ of $F(m)$, then the receipt of this message from i indicates that $s \in \hat{S}$. For example, a message which is in every $M(s)$ proves nothing and hence has no pure information content whatsoever. As we will see, even such messages can have significant equilibrium information content.

The receiver's payoff depends on his action and the true state of the world $s \in S$. Given a vector of messages, the receiver updates his beliefs

this individual is a homeowner and those in which the deed was forged.

[6] Implicit in this interpretation is the view that the receiver may be unaware of all implications of his prior knowledge.

possible $\underline{m}$.

We claim that if preferences are generic, then all equilibria in the induced game in state s have the receiver inferring state s. Suppose, contrary to the claim, that in some state s, there is an equilibrium in the induced game in which $\delta \neq s$ is inferred. First, suppose that no sender is "punished" in equilibrium. Hence sender 1 used a full report, the sender who contradicted him (if any), used a full report, etc. Therefore, the inference is some s' where some sender used the full report $m^*_{s'}$. This implies that $m^*_{s'} \in M(s)$, so $M(s') \subset M(s)$. By assumption, this implies that there is some i such that $s \succ_i s'$. At sender i's turn, there must have been some "claim on the table," established by an earlier full report, say $m^*_{s''}$. Since the true full report would have refuted this claim and been itself irrefutable, sender i could have guaranteed an inference of s. Hence $s \succ_i s'$ contradicts the assumption that this is an equilibrium.

Suppose then that some sender is punished in this equilibrium. Let sender k be the sender being punished. If $k = 1$, this means sender 1 did not use a full report. Otherwise, k is the first sender who contradicted the claim on the table but did not use a full report. Clearly, either way, sender k could have used m^*_s and guaranteed that the inference was s. Hence if this is an equilibrium, $\delta \succeq_k s$. But k is being punished, so that δ is the worst degenerate inference consistent with feasibility. Therefore, $\delta \preceq_k s$, implying $\delta \sim_k s$. By genericity, then, $\delta = s$. Hence there is no equilibrium in the induced game in state s in which the receiver's inference differs from s. Since there is necessarily a pure strategy equilibrium in the induced game in every state, we have a strong separating equilibrium.

It is easy to see that when preferences are not generic, the arguments above imply that in every equilibrium in the induced game for state s, either s is inferred or some sender k is punished and $\delta \sim_k s$. Clearly, there is an equilibrium in the induced game where we "break ties" by supposing that, if indifferent, senders avoid punishment. Hence even if preferences are not

APPENDIX

Proof of Proposition 1.

Suppose $s_1 \neq s_2$ but $M(s_1) = M(s_2)$. Clearly, if there is no separating equilibrium, there is no strong separating equilibrium. So suppose (σ, δ) is a separating equilibrium. By definition, then, (1) must hold for all i. Consider the strategies $\hat{\sigma}$, then, where $\hat{\sigma}_i^s = \sigma_i^s$ for every $s \neq s_2$ and $\hat{\sigma}_i^{s_2} = \sigma_i^{s_1}$. Since $M(s_1) = M(s_2)$, these strategies are feasible. Furthermore, the fact that (1) holds for every i for σ implies that the same is true of $\hat{\sigma}$. Clearly, though, $\delta(h_R(\hat{\sigma}, s_2)) = \delta(h_R(\sigma, s_1)) = s_1$. Hence (σ, δ) is not a strong separating equilibrium. ▌

Proof of Proposition 2.

Implied by Proposition 4.

Proof of Proposition 3.

Construct the receiver's inference rule as follows. Let m_s^* denote the full report for state s and, for a vector of messages $\underline{m} = (m_1, \ldots, m_n)$, let $F(\underline{m}) = \cap_{i=1}^n F(m_i)$. Suppose $\underline{m}$ is the vector of messages observed by the receiver. If $m_1 \neq m_s^*$ for any s, then the receiver's inference is that degenerate inference consistent with feasibility that sender 1 likes least. (If preferences are generic, this inference is unique. If they are not, choose any such inference.) If $m_1 = m_s^*$ and $s \in F(\underline{m})$, then the receiver infers s. Finally, suppose $m_1 = m_s^*$ but $s \notin F(\underline{m})$. Let k be the first sender for whom $m_k \notin M(s)$. Repeat the construction above with k in 1's place. More precisely, if $m_k \neq m_s^*$ for any s not ruled out by the messages of the first $k-1$ senders, then δ is the worst degenerate inference for sender k consistent with feasibility. If $m_k = m_{s'}^*$ for some $s' \in F(\underline{m})$, the receiver infers s'. If $m_k = m_{s'}^*$ but $s' \notin F(\underline{m})$, find the first $k' > k$ such that $\ldots$, etc. Clearly, this algorithm generates an inference for the receiver for each

and then chooses his action to maximize his expected utility. Let Δ be the set of probability distributions over S. The receiver's inference, then, is some $\delta \in \Delta$. The receiver's prior is $\delta^0 \in \Delta$, where we assume $\delta^0(s) > 0$ for all $s \in S$. When there is little risk of confusion, we will also use s to denote the probability distribution which puts probability 1 on state s. Such a probability distribution will be called a ***degenerate inference.***

Each sender i has preferences over the possible actions by the receiver, which in turn induce preferences over possible inferences by the receiver. In other words, sender i prefers the receiver's inference δ to δ' if he prefers the action the receiver takes given inference δ to the action taken given inference δ'.[7] Formally, the preference ordering of sender i is a complete, reflexive, and transitive ordering $\succeq_i$ over the set Δ. It is important to note that each sender i's preferences are independent of the true state.

We focus on a simple game, which we call an ***open forum***, in which each sender has one chance to send any one message he likes to the receiver. The sequence of senders is fixed and each sender, at his turn, observes the messages sent by previous senders. We number the senders in the order of their turns to speak. Given that each sender has only one chance to speak, the assumption that he can send only one message is without loss of generality. To see this, suppose that a sender could send either one or two messages from a set M. This is equivalent to allowing a sender to use one message from the set M' consisting of all single messages and pairs of

---

[7] An alternative interpretation of the model is that there are many symmetrically informed receivers who choose actions in some game after observing the senders' claims. If the equilibrium of this game is unique given each possible inference, we can restrict attention to the senders' preferences over inferences. With nonuniqueness, however, the senders' preferences would depend on the equilibrium among the receivers. Another difficulty with this interpretation is that we focus on how the receiver can obtain as much information as possible. With many receivers, depending on the nature of the game among them, more information could make all receivers worse off.

messages from the original set M.

To define senders' strategies, we first note that a sender i's information includes the true state s and the history of messages he has observed at the time he is called on to speak. Thus in state s, the history of messages observed by i is an element of $H_i(s) = [M(s)]^{i-1}$. A generic vector of messages will be denoted $\underline{m}$. Let Σ_i^s denote the set of functions $\sigma_i^s : H_i(s) \rightarrow M(s)$. A *strategy* for sender i is a collection of functions $\sigma_i = \{\sigma_i^s\}_{s \in S}$ where $\sigma_i^s \in \Sigma_i^s$ for each s. Since we focus on the existence of separating equilibria, we will not define mixed strategies.

Let the set of possible histories observed by the receiver be

$$H_R = \bigcup_{s \in S} [M(s)]^n .$$

Unlike senders, the receiver does not observe the state, only the messages sent. An *inference rule* is a function $\delta : H_R \rightarrow \Delta$, giving the receiver's beliefs as a function of the sequence of observed messages. A vector of strategies $\sigma = (\sigma_1, \ldots, \sigma_n)$ together with the state determines the particular history of messages, denoted $h_R(\sigma, s)$, observed by the receiver. Also, given any state s, any feasible messages $\underline{m}$ for the first $i-1$ senders, and any strategies for the remaining $n-i+1$, let $h_R(\underline{m}, \sigma_i^s, \ldots, \sigma_n^s)$ denote the history the receiver will observe.

An equilibrium specifies a vector of strategies σ for the senders and an inference rule δ for the receiver with the following properties. First, for each state s and each history $\underline{m}$, every sender i's strategy is optimal for him given the strategies of the senders after him and the receiver's inference rule. That is, for all i,

$$\delta[h_R(\underline{m}, \sigma_i^s, \sigma_{i+1}^s, \ldots, \sigma_n^s)] \succeq_i \delta[h_R(\underline{m}, m_i, \sigma_{i+1}^s, \ldots, \sigma_n^s)], \quad \forall m_i \in M(s), \underline{m} \in H_i(s), s \in S. \tag{1}$$

Second, the receiver's inference rule is consistent with the strategies of the senders. That is, he updates his beliefs using Bayes' Rule whenever

tions are independent of their information and (2) claims are only partially provable. With more than one speaker we found that a surprising amount of private information can be revealed even with only very limited provability. Indeed, conflicting preferences among the speakers is often sufficient for a robust form of full revelation which we call strong separation. Our results suggest that it is important to include communication with partial provability in models of real world economic, legal, and political processes because even minimal amounts of provability may radically affect predicted outcomes.

Many open questions and issues remain for future work on partial provability. For example, Proposition 5 and Examples 2 and 3 clearly indicate that weaker sufficient conditions for strong separation are possible. Also, it would be interesting to characterize the structure of pooling equilibria. Other interesting conceptual issues include: What happens if the receiver has the ability to ask questions? What effect would this have on the amount of information he can obtain? Can we characterize the "welfare loss" caused by information which is not revealed?

The analysis presented here was at a fairly general level. Exploiting the additional structure on preferences and message structures available in specific applications will yield sharper insights into the role of communication in these activities. For example, what rules of evidence and order of presentation maximize the incentive of litigants to reveal their information? What debate rules best encourage candidates to truthfully state their qualifications and positions?

The intuition of this example is straightforward. In state s_2, the two senders would rather trick the receiver into inferring s_1. For a separating equilibrium to exist, the receiver must infer s_1 if he sees m_1 from both agents. Since this is feasible in state s_2, there is a way for them to trick him. In the sequential game, sender 1 can "offer" this deception to sender 2 by sending message m_1. Sender 2 will follow his lead. In the simultaneous game, though, they cannot collude this way. If sender 2 expects sender 1 to use m_1 and so uses this message himself, sender 1 will actually "doublecross" him and send m_3 to get the receiver to infer s_3.

The key reason why the open forum is easier to analyze than the simultaneous game is that the induced game, given a state and the receiver's inference rule, is a finite game of perfect information. This fact implies that pure strategy equilibria always exist in this induced game, no matter what the receiver's inference rule is. Hence if the inference rule can be chosen to insure that there is no equilibrium in the induced game where the receiver infers incorrectly, there must be a strong separating equilibrium. (This fact also explains why it is sometimes just as easy to prove that a strong separating equilibrium exists as it is to prove that a separating equilibrium exists.) Our open forum inference rules when applied to the simultaneous game do guarantee that there is no equilibrium in the induced game in which the receiver is misled. However, in the simultaneous game, the absence of an equilibrium in which the receiver is misled does not guarantee the existence of an equilibrium in which he infers the truth. Because the induced game in the simultaneous mechanism is not a game of perfect information, there may be no pure strategy equilibrium.

V. Conclusion.

This paper analyzed the ability of an uninformed decision maker to elicit private information from interested parties. In particular, the decision maker in our model must overcome two important obstacles to informative communication: (1) speakers' preferences over possible decision maker's ac-

the history he observes has nonzero probability under the strategies σ. Formally, if, given a history $h_R \in H_R$,

$$S^*(h_R,\sigma) = \{s \in S \mid h_R = h_R(\sigma,s)\}$$

is nonempty, then

$$\delta(h_R)(s) = \begin{cases} \delta^0(s)/K, & \text{if } s \in S^*(h_R,\sigma); \\ 0, & \text{otherwise;} \end{cases} \tag{2}$$

where $K = \sum_{s \in S^*(h_R,\sigma)} \delta^0(s)$. Third, off the equilibrium path, we only require that his inferences are ***consistent with feasibility*** in the sense that his beliefs give zero probability to any state in which the messages he has observed are infeasible. In other words, if $S^*(h_R,\sigma)$ is empty, then $\delta(h_R)$ is some belief such that

$$h_R \notin [M(s)]^n \quad \Rightarrow \quad \delta(h_R)(s) = 0.$$

This is essentially Fudenberg and Tirole's [1991] perfect Bayesian equilibrium. One may wish to impose other requirements on the receiver's beliefs off the equilibrium path. For example, sequential equilibrium would require the receiver to view a unilateral deviation from σ as infinitely more likely than two or more senders deviating.[8] While we do not require this, all our results hold with this additional restriction.

A *separating equilibrium* is an equilibrium (σ,δ) in which $\delta(h_R(\sigma,s)) = s$ for all s. In a separating equilibrium, then, the receiver learns the state exactly.

An alternative way to view equilibria is useful. Imagine that the receiver can announce the inference rule he will follow. He cannot commit himself to drawing "irrational" inferences in the sense that his inference rule must be part of an equilibrium. However, he can commit himself to following any inference rule satisfying this criteria. Given δ and s, the

[8] This requirement plays a prominent role in Kreps and Ramey [1987] for example.

senders, in effect, play a game among themselves. It is not hard to see that if (σ, δ) is an equilibrium, then the senders' strategies must also form an equilibrium of this "subgame" for each s. (This is nothing more than restating requirement (1) above.)

Formally, given δ and s, the ***induced game*** is the game of perfect information where sender i's strategy set is Σ_i^s and where his preferences over histories of play $\underline{m} \in [M(s)]^n$ are those induced by δ. That is, in the induced game, sender i weakly prefers $\underline{m}$ to $\underline{m}'$ iff $\delta(\underline{m}) \succeq_i \delta(\underline{m}')$.

This intuition is useful because it suggests a stronger criterion for equilibria. Fix an equilibrium (σ, δ). Suppose that there are equilibria other than σ in the induced game given δ. If some of these equilibria lead the receiver to an incorrect inference, then the receiver must know which equilibrium the senders are playing in order to know that his inferences are correct. Alternatively, if all equilibria in the induced game make his inference rule correct, he does not have this problem. In this case, we will say that the equilibrium is ***strong***. This requirement does not say that there are unique equilibrium ***strategies*** in the induced game, only that there is a unique equilibrium ***outcome***.

More formally, an equilibrium (σ, δ) is strong if for every σ' satisfying (1) for every i, we have $\delta(h_R(\sigma', s)) = \delta(h_R(\sigma, s))$ for all s. Hence, in particular, a separating equilibrium is strong if for every σ' satisfying (1) for every i, we have $\delta(h_R(\sigma', s)) = s$ for all s.

Remark. The reader familiar with the literature on full implementation may find a comparison useful. In this literature (see Maskin [1977], Moore and Repullo [1988], or Palfrey and Srivastava [1991]), one seeks a game form, say (M, O), which fully implements a social choice correspondence. More precisely, $M = \prod_i M_i$ where M_i is the message set for agent i. O is an outcome function, mapping M into some outcome space. The agents all have preferences over outcomes where these preferences depend on the

Note, also, that m_3 from sender 1 and m_1 from sender 2 is an equilibrium of the induced game in this state. Finally, consider state s_2. It is easy to see that both senders using message m_2 is an equilibrium in the induced game in this state — since sender i's opponent is proving that s_2 is true, sender i may as well do so too. Also, no other inference can be made in an equilibrium. To see this, note that there is no equilibrium in the induced game where the receiver infers s_1. For this to happen, both senders must be using m_1, in which case sender 1 prefers deviating to m_3. Also, it cannot be true that s_3 is inferred, since sender 2 would deviate to m_2. Hence we have a strong separating equilibrium.

By contrast, there is no separating equilibrium in the open forum, strong or otherwise. To see this, suppose not. Clearly, the separating equilibrium must have the receiver infer s_1 if both senders use m_1 — no other messages are possible in s_1. Given this, sender 1 must use m_3 in state s_3. If he uses m_1, sender 2 could reply with m_1 and achieve his favorite inference of s_1. Also, notice that s_3 is sender 2's least favorite inference. Hence it must be true that if sender 1 uses m_3, the receiver will infer s_3 given any message by sender 2 which does not disprove s_3. If this were not the case, then sender 2 could avoid the inference s_3 even when s_3 is true. Finally, for the receiver's inference rule to be consistent with feasibility, he must infer s_2 if either sender uses message m_2.

So consider state s_2. If sender 1 uses message m_2, he proves that s_2 is true. If he uses message m_3, sender 2 has two choices: use message m_2 to prove that the state is s_2 or do not disprove s_3, leading to inference s_3. Clearly, he prefers the former. Finally, suppose sender 1 uses m_1. Since sender 2 can respond with m_1 to get his favorite inference of s_1, he will certainly do so. Therefore, sender 1 can generate the inference s_2 (by sending message m_2 or m_3) or the inference s_1 by sending m_1. He prefers the latter, so the receiver does not end up inferring correctly. Hence there is no separating equilibrium.

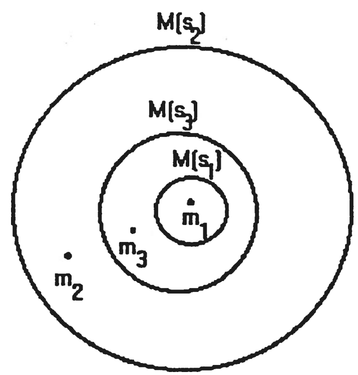

Figure 5.

state of the world. A social choice correspondence is fully implemented by this game form if, for each state, the set of equilibrium outcomes equals the social choice set for the state. By contrast, in our model the set of feasible messages, not preferences, vary with the state. Our notion of strong separation is very similar to the notion of full implementation. With strong separation, given a state and the inference rule, all equilibria have the same outcome. With full implementation, all equilibria are "acceptable" in the sense that their outcomes satisfy the social choice rule. In our model, we view the receiver as a player in the game who cannot commit himself to choosing suboptimally given some inference. In the full implementation literature, the only "receiver" is the mythical social planner who can commit to any (feasible) outcome out of equilibrium.

III. Results.

As discussed in the introduction, complete provability models assume that every true statement is provable. That is, for every $\hat{S} \subseteq S$, there is a message $m_{\hat{S}}$ such that $F(m_{\hat{S}}) = \hat{S}$. As far as the results on separation in Milgrom [1981] and Grossman [1981] are concerned, the key aspect of this assumption is that for every $s \in S$, there is a message m_s which proves s — that is, $F(m_s) = \{s\}$. This assumption can be thought of as having two components. First, if s is the true state, then for any other state, s', there is a way to disprove s'. That is,

(M1) $$s \neq s' \;\Rightarrow\; M(s) \not\subseteq M(s').$$

Second, there is a single message available in state s which proves by itself what all messages available in s prove jointly. That is, for every s, there exists m_s^* such that

(M2) $$F(m_s^*) = \bigcap_{m \in M(s)} F(m).$$

In other words, if a message $m_1 \in M(s)$ disproves every $s' \in \hat{S}_1$ and a message $m_2 \in M(s)$ disproves every $s' \in \hat{S}_2$, then there is a message

$m_3 \in M(s)$ which disproves (at least) every $s' \in \hat{S}_1 \cup \hat{S}_2$. When assuming (M2), we will refer to the message m_s^* as the *full report* for state s.

It is important to note that, in general, the full report for state s does not prove that the true state is s. In particular, if there is another state s' such that $M(s) \subseteq M(s')$, then the full report for state s is also feasible in s' and hence does not disprove s'. Note, though, that this is the only kind of s' that the full report for s would not rule out. Since (M1) eliminates this possibility, these two conditions together imply that the full report for s does prove that s is the true state.

Clearly, both (M1) and (M2) are quite strong. As discussed in the introduction, it may well be impossible to provably rule out some possibility even when it is false. In such a situation, (M1) does not hold. Also, in situations where complete proof requires more time or space than is available, (M2) will not hold. Hence results with partial provability — *i.e.*, with one or both of (M1) and (M2) relaxed — would be desirable. Previous models of partial provability relax either assumption (M1) or assumption (M2). Both the "not" message structure of Example 1 and the "any–k–signals" structure of Milgrom [1981] and Fishman and Hagerty [1990][9] satisfy (M1) but not (M2). The "not–less–than" message structure of Okuno–Fujiwara, Postlewaite, and Suzumura [1990][10] satisfies (M2) but not (M1). One purpose of this paper is to explore more systematically the effect of relaxing these assumptions.

---

9 In an example, Milgrom [1981] assumes that the sender observes N signals. The set of feasible messages is the set of truthful disclosures of at most k of these signals. In Fishman and Hagerty [1990], each signal takes on one of two values, "high" or "low," and $k = 1$.

10 They assume that a sender observes a signal from a finite, ordered set. The set of messages is the set of truthful lower bounds for the signal.

Since a "not" message eliminates exactly one other state, we see that any assignment of messages to states covers at most $2(\ell - 1)$ of these pairs. Thus the number of pairs to cover exceeds the number that can be covered whenever $\ell > 4$. In other words, with two senders, strictly conflicting generic preferences, and the "not" messages, the simultaneous game does not have a strong separating equilibrium if $\ell > 4$. On the other hand, as seen in Example 1, under the same conditions, the open forum has a strong separating equilibrium for any ℓ. Hence the simultaneous mechanism does not dominate the open forum in terms of strong separation.

The following example shows that the open forum does not dominate either.

Example 5.

There are two senders and three states, s_1 through s_3. The message sets are shown in Figure 5. The preferences over degenerate inferences are $s_3 \succ_1 s_1 \succ_1 s_2$ and $s_1 \succ_2 s_2 \succ_2 s_3$. Preferences over nondegenerate inferences can be chosen in any way such that each sender's favorite inference and least favorite inference are degenerate.

To see that the simultaneous mechanism has a strong separating equilibrium, suppose the receiver infers as follows. If both senders use the message m_1, he infers s_1. If either sender uses the message m_2, he infers s_2. (Since m_2 proves that s_2 is the true state, this is necessary for the inference rule to be consistent with feasibility.) Finally, for any other pair of messages, he infers s_3. Consider the game this induces among the senders. Obviously, in state s_1, only message m_1 is feasible, so, trivially, the unique equilibrium in the induced game in this state has both senders using m_1. In state s_3, it is easy to see that all equilibria have s_3 inferred. It is impossible for the receiver to infer s_2 as the message m_2 is not feasible. Hence the only way he could infer something different from s_3 is if both senders use m_1. Clearly, sender 1 would deviate to m_3 to achieve the inference s_3.

Proposition 8. *When $n = 2$ and senders 1 and 2 have generic, conflicting preferences, there is a δ such that (σ, δ) is a strong separating equilibrium if for all s, s' such that $s \succ_1 s'$, either*

$$\sigma_1(s) \notin M(s')$$

or

$$\sigma_2(s') \notin M(s).$$

Furthermore, if preferences are strictly conflicting in the sense that $\delta \succ_1 \delta'$ iff $\delta' \succ_2 \delta$, then there is no separating equilibrium (strong or otherwise) unless such a σ exists.

The intuition of this result is very simple. Suppose that $s \succ_1 s'$ so $s' \succ_2 s$ and that the condition of the theorem does not hold. What should the receiver infer if he gets the message $\sigma_1(s)$ from sender 1 and the message $\sigma_2(s')$ from sender 2? If he infers that the state is s, then sender 1 deviates from the equilibrium in state s'. Similarly, if he infers it is s', sender 2 deviates from the equilibrium in state s. Suppose, then, that he draws some completely different inference, say δ. If $\delta \succ_1 s'$, sender 1 still deviates in state s'. Hence we must have $\delta \preceq_1 s'$. But by strictly conflicting preferences, whether δ is degenerate or nondegenerate, this implies $\delta \succeq_2 s' \succ_2 s$. Thus sender 2 deviates in state s. In short, if the condition of the theorem does not hold, we cannot simultaneously prevent sender 1 from deviating in state s and sender 2 from deviating in state s'.

To see one implication of Proposition 8, consider the "not" messages of Example 1. Note that Proposition 8 requires that, for every pair of states, either sender 1 or sender 2 must "cover" that pair in the sense that his message in his preferred state is not feasible in the less preferred state. With ℓ states, the number of pairs to cover is $\ell(\ell - 1)/2$. However, the number of messages we can assign is only one for each sender for each state, or 2ℓ. Notice further that only $2(\ell - 1)$ of these can cover a pair, since a sender's message in his least favorite state cannot cover a pair.

In relaxing these assumptions on the message sets, we do not wish to introduce trivial barriers to communication. In particular, when (M2) is dropped, it is possible to have fewer messages than states. In such cases, there is no way for a sender to unambiguously claim the truth in every state. To avoid this pathology, we assume throughout that a sender can always include in his message a claim of any state that the message does not disprove. More formally, for every message m,

$$\#\{m' \mid F(m') = F(m)\} \geq \#F(m).$$

We refer to this as the cardinality condition. It guarantees that the receiver can treat different messages with the same pure information content as claiming different states with the same proof. In describing our results and examples, we will be less formal, assuming that senders can include a "cheap talk" claim of a state with their message.

Strong separation will require some less innocuous restrictions. Consider the following substantial weakening of (M1):

(M3) $$s \neq s' \quad \Rightarrow \quad M(s) \neq M(s').$$

Given that the preferences of the senders are independent of the state, our first result is perhaps unsurprising.

Proposition 1. *(M3) is necessary for the existence of a strong separating equilibrium.*

What is more surprising is that (M3) is *not* necessary for separation *per se*. As a trivial example, suppose one of the senders is completely indifferent over all possible inferences by the receiver. Clearly, there is an equilibrium in which he tells the truth regardless of whether (M3) holds. However, this need not be a strong separating equilibrium. In particular, holding the receiver's inference rule fixed, this sender has many other best replies, some of which may not lead to the correct inference. It is precisely this kind of multiplicity of equilibria in the induced game that leads us to focus on strong separating equilibria. (It is not difficult to construct less trivial examples.)

As discussed in the introduction, our goal is to replace strong assumptions on provability with weak assumptions on preferences. The main such assumption is the following.

Definition. *Senders have* conflicting preferences *if $s \neq s'$ implies that there is some i such that $s \succ_i s'$.*

With only two senders, this assumption implies that their preferences over degenerate beliefs disagree on every comparison. With more than two, it only says given a pair of degenerate beliefs, there is some pair of senders who disagree.[11] In particular, this assumption says nothing about comparisons between nondegenerate beliefs. Another assumption will also be used for some of our results.

Definition. *Senders have* generic preferences *if for all i and for all $s \neq s'$, we have $s \not\sim_i s'$.*

That is, with generic preferences, no sender is indifferent between any pair of degenerate inferences.

Our first sufficiency result generalizes the "not" example in the introduction. Recall that the key to that equilibrium was that sender 1's message included a claim of a particular state, where this claim was believed unless refuted by sender 2. Given the "not" message structure, sender 2 was able to refute any false claim by 1. Their conflicting preferences guaranteed that

---

[11] Milgrom and Roberts [1986] use a similar assumption in showing that the equilibrium outcome is unique.

An equilibrium is a vector of strategies σ and an inference rule δ such that (1) the senders are using best replies and (2) δ is consistent with Bayes' Rule when possible and consistent with feasibility always. The definitions of a separating equilibrium and a strong separating equilibrium carry over in the obvious manner.

One striking fact about simultaneous games is that separating equilibria exist under extremely weak assumptions — even when $M(s) = M$ for all s! However, these equilibria are not strong. Thus the case for focusing on strong separation is even more compelling for simultaneous games than for the open forum. As an example, suppose there are ℓ states and ℓ distinct messages $m_1, \ldots, m_\ell$ which are feasible in every state. Since there are many ways to prove nothing, this assumption seems quite reasonable. Then if there are at least three senders, regardless of what other messages are feasible and regardless of the preferences of the senders, a separating equilibrium exists in the simultaneous game. To see this, suppose that each sender i in state s_k sends message m_k. The receiver's inference rule is to infer s_k if at least $n-1$ senders send message m_k. This is a separating equilibrium since any one sender cannot change the receiver's inference by deviating. However, in general, this is not a strong equilibrium. For example, suppose all the senders prefer the inference s_1 to any other inference. Then, given this inference rule, it is also an equilibrium in the induced game for all senders to send message m_1 in every state because no sender would want to disprove s_1. More generally, it is easy to see that (M3) is necessary for the existence of strong separating equilibria of the simultaneous game or, for that matter, of *any* game.

Our main result of this section is that neither the simultaneous mechanism nor the open forum dominates the other as far as existence of strong separating equilibria is concerned. To show this, we begin with a simple characterization of when a simultaneous game has strong separating equilibria for the special case of two senders and generic, conflicting preferences.

any state such that $s_1 \succeq_1 s'$ for all s'. Because of the preferred message property, a necessary condition for a strong separating equilibrium is that there exists $m_1 \in M(s_1)$ which is not in $M(s')$ for any $s' \neq s_1$. Let s_2 denote any state such that $s_2 \succeq_1 s'$ for all $s' \neq s_1$. First, suppose $s_1 \succ_1 s_2$. Then there must be a message $m_2 \in M(s_2)$ which is not in $M(s')$ for any $s' \notin \{s_1, s_2\}$. Alternatively, suppose $s_1 \sim_1 s_2$. Then there must be a message $m_2 \in M(s_2)$ which is not in $M(s')$ for any $s' \neq s_2$, etc. Condition (4) is exactly what this construction requires — that is, it is necessary and sufficient for the existence of a strategy for sender 1 which could be part of a strong separating equilibrium.

Condition (4) is not sufficient because it does not guarantee that we can construct an inference rule which generates strong separation. To see this intuitively, suppose there are only two states, s_1 and s_2, where $M(s_1) \cap M(s_2)$ is not empty but neither set is a subset of the other. Clearly, regardless of the preferences of sender 1, condition (4) holds. Suppose, though, that sender 1 is completely indifferent over all inferences by the receiver. Then no matter what inference rule the receiver follows, sender 1 will necessarily have a best reply for which the receiver does not infer the true state exactly. However, if preferences are generic, this kind of problem is eliminated.

IV. Simultaneous Games.

In this section, we briefly discuss separation in games where the senders send messages to the receiver simultaneously — or, equivalently, without having an opportunity to see the messages being sent by any other sender. Our goal is to contrast the possibilities for separation in the open forum versus this simultaneous game.

The definitions of strategies are simpler here. A strategy for sender i is simply a function $\sigma_i : S \to M$ such that $\sigma_i(s) \in M(s)$ for all s. As before, the receiver's inference rule maps vectors of messages into beliefs.

2 had an incentive to refute false claims.

A simple extension of this idea to n senders is the following. Our cardinality condition allows us to assume that senders include a "cheap talk" claim of a state along with their message. Think of sender 1's claim as putting this state "on the table." This claim remains on the table until refuted by a subsequent sender who replaces it with a claim of his own. The claim which is on the table at the end of the game is then believed by the receiver. We will refer to this kind of inference rule as a BUR rule, short for "believe unless refuted."

As in the "not" example, there are two requirements for achieving strong separation with a BUR rule: senders must have (1) the ability and (2) the incentive to refute any false claim. The former is precisely what assumption (M1) guarantees. A relatively straightforward way to guarantee the latter is to assume conflicting preferences. Thus we have the following proposition.

Proposition 2. *If (M1) and conflicting preferences hold, there is a strong separating equilibrium.*

To see the intuition behind this result, suppose the receiver follows the BUR rule. Suppose there is an equilibrium in the induced game in which the receiver infers a state s' other than the truth. Clearly, the true state is never put on the table in this equilibrium. By conflicting preferences, there is a sender who prefers that the receiver infer the truth. By (M1), no matter what state is on the table at this sender's turn, be it s' or any other, he can refute it and claim the true state. Since no sender can refute the truth, this will be what the receiver infers. Hence there are no equilibria in which the receiver infers incorrectly. Since the induced game is a finite game of perfect information, a pure strategy equilibrium exists for each state. Hence there is an equilibrium in which the receiver infers correctly. Therefore, a strong separating equilibrium exists.

Does the BUR rule work if (M1) is relaxed? The difficulty is illustrated in Figure 1 below. Suppose that the receiver uses the BUR rule and that sender 1 uses m_2 to claim s_1 when the true state is s_2. Since $M(s_2) \subset M(s_1)$, it is impossible to refute the claim s_1. Hence the receiver would incorrectly infer that sender 1 told the truth. However, a simple modification of the naive BUR rule avoids this problem. If the receiver only "allows" sender 1 to claim s_1 if his message rules out s_2, then he can safely believe this claim unless it is subsequently refuted. More precisely, if sender 1 claims s_1 without ruling out s_2, the receiver treats this as a claim of s_2. More formally, recall that that the association of messages with claims is part of the receiver's inference rule. Hence he can use an inference rule which only associates certain claims with evidence satisfying certain criteria.

When (M2) holds, a natural restriction is that a claim of s must be accompanied by the full report for s. We call this the "believe full reports unless refuted" or BFUR rule. The BFUR rule works with conflicting preferences, but, as the following proposition shows, an even weaker preference condition is sufficient.

Proposition 3. *If (M2) and (M3) hold and if*

(3) $$M(s) \subset M(s') \;\Rightarrow\; \exists i \text{ such that } s' \succ_i s,$$

then there is a separating equilibrium. If, in addition, preferences are generic, a strong separating equilibrum exists.

To see why this works, suppose sender i uses the full report m_s^* for state s. By definition, this proves that every message $m \in M(s)$ is feasible in the true state. It does not prove that other messages are infeasible. Hence if $s' \neq s$ is the true state, it must satisfy $M(s) \subseteq M(s')$. By (M3), then, $M(s) \subset M(s')$, so that it is possible to refute sender i's claim of s. The preference assumption (3) guarantees that some sender has an incentive to do so. The genericity assumption is only used to ensure that senders will use

Definition. *The preference order $\succ_i$ satisfies* betweenness *if for all non-degenerate δ, we have $\underline{\delta} \prec_i \delta \prec_i \bar{\delta}$ where $\underline{\delta}$ is the worst and $\bar{\delta}$ the best degenerate inference for sender i consistent with s in the support of δ.*

Notice that betweenness is weaker than convexity of preferences.

Proposition 7. *When there is only one sender, if (4) holds and $\succ_1$ satisfies genericity and betweenness, then all equilibria, in pure or mixed strategies, strong or not, are separating.*

It is straightforward to characterize strong equilibria with one sender because every such equilibrium, separating or not, has what we call the ***preferred message property***. More precisely, the sender "ranks" the equilibrium messages according to the inferences they generate and sends the highest ranked feasible message in each state (where two or more messages may have the same rank). Furthermore, for every s, there is a unique highest ranked feasible message. To see this, fix a strong equilibrium and consider the set of messages the sender uses in equilibrium. Each of these messages leads to a different inference by the receiver. Let $m_1, \ldots, m_k$ denote the messages which lead to the best (for sender 1) of these inferences. Clearly, then, if one of these messages is feasible, it is sender 1's choice. Since the equilibrium is strong, sender 1 cannot have multiple (feasible) best replies to the receiver's inference rule. That is, $F(m_i) \cap F(m_j) = \emptyset$ for all $i \neq j$. Hence for every $s \in \cup_{i=1}^k F(m_i)$, sender 1 sends the unique message from the set $\{m_1, \ldots, m_k\}$ which is feasible in s. Let $m'_1, \ldots, m'_{k'}$ be the messages leading to the second favorite. If none of $m_1, \ldots, m_k$ is feasible but one of $m'_1, \ldots, m'_{k'}$ is, then sender 1 uses one of these messages, etc. It is worth noting that this property does not carry over to the multiple sender case even as a description of best replies.

Because all one–sender strong equilibria have this property, it is easy to see that condition (4) is necessary for strong separation. Let s_1 denote

On the other hand, suppose sender 2 has the same preferences as sender 1. In this case, there is a strong separating equilibrium. Since sender 2 wants s_1 to be inferred, he will always say the opposite of the message used by sender 1 in this state, proving that s_1 is true.

Proposition 5 does *not* extend to the case of only one sender. This fact is a consequence of Proposition 6 which is a generalization of a result in Okuno–Fujiwara, Postlewaite, and Suzumura [1990].[14] For any state s and any sender i, let

$$L_i(s) = \{s' \neq s \mid s' \preceq_i s\}.$$

Proposition 6. *When there is only one sender, a separating equilibrium exists if for every s,*

$$M(s) \setminus \left[\bigcup_{s' \in L_1(s)} M(s') \right] \neq \emptyset. \tag{4}$$

If preferences are generic, (4) is necessary and sufficient for the existence of a strong separating equilibrium.

Clearly, the requirement that there be some ordering of the states such that (4) holds is much stronger than (M3). In particular, if there are preferences for which (4) holds, then for any best inference s under these preferences, there must be a message in $M(s)$ which rules out *every* other state.

Under a weak additional condition, we can say still more about the separating equilibria of the one–sender case.

[14] They consider multiple asymmetrically informed senders. However, the information known to sender i is independent of that known to j. Furthermore, their preference assumptions imply that what sender i wishes to signal is independent of what the other senders are signaling. Hence, in effect, they have several one–sender games which operate side–by–side.

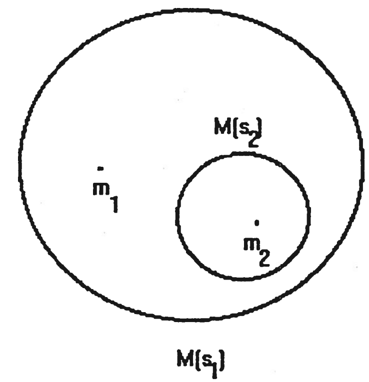

Figure 1.

full reports in every equilibrium. In particular, if the receiver can effectively ban the use of messages which are not full reports, then genericity is not needed for the existence of a strong separating equilibrium.

The existence of full reports, (M2), is a much stronger condition than is needed for a modified BUR rule to work. Suppose instead that the message sets have the property that in claiming s, a sender can simultaneously rule out every s' such that $M(s') \subseteq M(s)$. By Proposition 1, strong separation requires $M(s') \neq M(s)$, so we can restate this condition to say that in claiming s, a sender is able to simultaneously rule out s' if its message set is a strict subset of $M(s)$. For obvious reasons, it is only necessary to rule out such an s' if no sender has already ruled it out. We will call a claim of s *trustworthy* if it is made using a message which rules out every s' not ruled out by previous messages such that $M(s') \subset M(s)$. When will a BTUR ("believe every trustworthy claim unless refuted") rule work? Clearly, it must be possible to claim the truth in a trustworthy fashion. Furthermore, for every sender except the first, this must be done while refuting some existing claim. This intuition suggests the following condition:

$$\text{(M4)} \qquad s \notin \hat{S}(s') \;\Rightarrow\; M(s) \setminus \left[\bigcup_{s'' \in \hat{S}(s)} M(s'') \right] \not\subseteq M(s')$$

where, for any state s,

$$\hat{S}(s) = \{s'' \mid M(s'') \subset M(s)\}.$$

To understand this condition, suppose that s' is the claim on the table and s is true. Since a trustworthy claim of s' must include refutation of any possible state s'' such that $M(s'') \subset M(s')$, it must be true that $s \notin \hat{S}(s')$. Condition (M4) guarantees that a sender can make a claim s while ruling out *every* $s'' \in \hat{S}(s)$ (regardless of whether some of these s'' were previously ruled out) and simultaneously refuting s'. Clearly, such a claim of s is trustworthy.[12] We then have the following proposition.

12 One can substantially relax (M4) by changing the definition of a trustworthy claim.

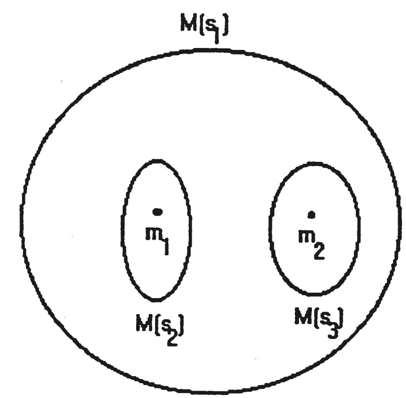

Figure 4.

for the result: the remaining senders' preferences are completely unconstrained. Also, the only relevant aspect of preferences are how degenerate inferences are ranked. Preferences over other inferences can be chosen arbitrarily.

An implication of Proposition 5 is that necessary and sufficient conditions for the existence of strong separating equilibria will not break neatly into a set of conditions on messages and a set of conditions on preferences. Since there are no conditions on the messages other than (M3) which are necessary independently of preferences, any necessary and sufficient conditions must interrelate the two aspects of the problem.

By Proposition 1, (M3) is necessary for the existence of a strong separating equilibrium. Hence we can paraphrase Proposition 5 as saying that if the message technology permits a strong separating equilibrium, then it permits one when we assume conflicting preferences. This may appear to confirm the old intuition that the best way to elicit information is to choose senders whose interests are in direct conflict. This approach seems to create the most intensive possible "informational competition" and hence to generate the most information. However, as the next example shows, this intuition is not correct. Certain conflicts of preferences among the senders are worse for the receiver than commonality of interests. Proposition 5 only guarantees that the *right* conflicting preferences generate strong separation.

Example 4.

There are three states, s_1 through s_3, and two senders with preferences $s_1 \succ_1 s_2 \succ_1 s_3$ and the reverse ordering for sender 2. Hence conflicting preferences hold. The message sets are shown in Figure 4. Clearly, there is no separating equilibrium. If there were one, then obviously the receiver would have to infer s_2 if he sees m_1 from both senders and s_3 if he sees m_2 from both. But then in state s_1, sender 2 would repeat whatever message sender 1 uses to avoid the inference s_1.

Proposition 4. *If (M3), (M4), and conflicting preferences hold, there is a strong separating equilibrium.*

Proposition 2 is a direct implication of this result as (M1) implies both (M3) and (M4). Because of its different preference assumptions, Proposition 3 is not a corollary. However, the message assumption there is strictly stronger: (M2) also implies (M4).[13]

Is (M4) a necessary condition for strong separation? Clearly, it would be not be easy to relax this condition given the BTUR rule. However, there are many other types of inference rules which achieve strong separation under weaker message conditions. The next two examples illustrate alternative approaches. The first example works by strengthening the requirements for a claim to be "believed unless refuted," while the latter works by having the receiver disbelieve some nonrefuted claims.

Example 2.

There are three states, s_1, s_2, and s_3, and two senders with preferences $s_1 \succ_1 s_2 \succ_1 s_3$ and $s_3 \succ_2 s_1 \succ_2 s_2$. The message sets are shown in Figure 2. (M1) does not hold as $M(s_1) \subset M(s_2)$. Also, (M2) does not hold. There is no full report for state s_2 since m_2 and m_3 together prove that s_2 is true, though there is no one message which proves this. Even the weaker (M4) does not hold as $M(s_2) \setminus M(s_1) = \{m_3\} \subset M(s_3)$. Finally, neither conflicting preferences nor the assumption (3) used in Proposition 3 are satisfied. Only (M3) holds.

Despite this, there is a strong separating equilibrium. The BTUR rule

In particular, it is only necessary for a sender claiming s to rule out any state s'' not already ruled out such that $M(s'') \subset M(s)$ and $s \succ_1 s''$.

[13] To see this, note that the full report for s always rules out every s'' such that $M(s'') \subset M(s)$ and is only contained in $M(s')$ if $s \in \hat{S}(s')$.

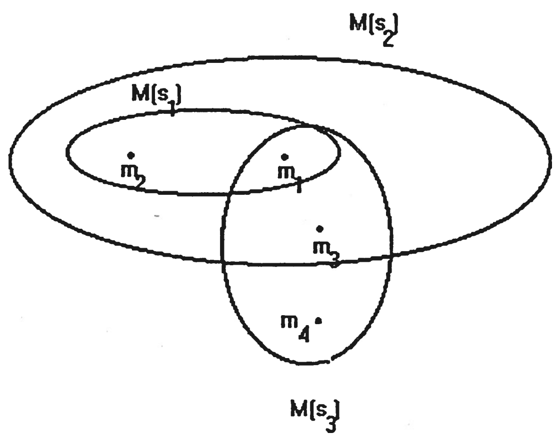

Figure 2.

the message sets other than (M3) which are necessary for strong separation independently of the preferences. Put differently, with the "right" preferences, strong separation is possible as long as (M3) holds. Furthermore, this statement is true even with some rather stringent restrictions on the preferences one can "choose."

Proposition 5. *If (M3) holds and there are at least two senders, then there exist preferences $\succ_1$ and $\succ_2$ for the first two senders satisfying genericity and conflicting preferences such that, regardless of the preferences of the remaining senders, a strong separating equilibrium exists.*

To see the intuition of this result, suppose there are exactly two senders. Choose any generic preferences for sender 1 such that $M(s) \subset M(s')$ implies $s \succ_1 s'$. Choose any preferences for sender 2 satisfying conflicting preferences. (It is not hard to show that this is possible.) Consider the following inference rule. For any state s, sender 1 can use any message $m \in M(s)$ to claim s. If sender 2 does not refute 1's claim, the receiver infers that 1 told the truth; otherwise, the receiver infers the worst state for sender 1 consistent with the feasibility of the messages he has seen. All equilibria in the induced game given this inference rule are separating. If sender 1 claims s, there are three possibilities. First, s may be true, in which case sender 2 cannot refute it so s will be inferred. Second, s may be false and the true state, s', satisfies $M(s') \not\subset M(s)$. In this case, sender 2 can refute sender 1's lie, thereby generating an inference better for himself — and hence worse for 1 — than s'. In this case, sender 1 should have claimed the truth instead. Finally, it is possible that s is false and that the true state, s', satisfies $M(s') \subset M(s)$. In this case, sender 2 cannot refute 1's claim, so s will be inferred. However, by construction, $s' \succ_1 s$, so sender 1 would have been better off claiming the truth. More colloquially, this equilibrium works because the only lies sender 2 could not refute are lies sender 1 would not want to tell.

As this explanation suggests, only the first two senders are important table. If he could avoid confirming sender 2's implicit challenge and still obtain the inference s_1, he would certainly do so. Thus, in equilibrium, we could have sender 1 claiming s_1, sender 2 claiming s_2 without refuting s_1, and sender 3 not refuting either claim leading the receiver to infer s_2.

A common feature of Examples 2 and 3 is that the following relaxation of (M4) is satisfied:

$$\text{(M5)} \qquad M(s) \setminus \left[\bigcup_{s' \in \hat{S}(s)} M(s') \right] \neq \emptyset.$$

If (M5) holds, it is always possible to claim s in a trustworthy fashion. However, as seen in Example 3, it may not be possible to do so while simultaneously refuting some prior claim on the table. Hence (M5) guarantees that either an implicit or explicit challenge with a claim of the truth is always possible. A natural conjecture is that (M5) together with (M3) and conflicting preferences implies the existence of a strong separating equilibrium. In fact, this is true for the two and three sender cases. The proof for the two–sender case is straightforward: any false claim by sender 1 which is better for him than the truth is worse for sender 2 than the truth. Since trustworthy claims, if false, can always be refuted, sender 2 can refute such a claim and clearly has the incentive to do so.

The three–sender case, on the other hand, is quite complex. As the example above indicates, when (M5) is satisfied but (M4) is not, some device like implicit challenges must be introduced. Formulating an inference rule which rewards sender 2 for challenging false claims by sender 1 but which does not allow senders 1 and 2 to jointly make a false claim that 3 cannot block is quite complex. The proof that such an inference rule exists is available from the authors on request. Unfortunately, it does not appear straightforward to generalize this rule to the n–sender case.

Even (M5), however, is not necessary — not even if we assume conflicting preferences. As the next result shows, there are *no* conditions on

only requires that claims of s_2 must be made using m_3 as this is necessary to rule out s_1. Consider an alternative inference rule which requires a claim of s_1 to be made using m_1 — that is, claiming s_1 with m_2 is not allowed. To see that this generates a strong separating equilibrium, suppose sender 1 makes a false claim. If the true state is s_3, then sender 2 can obtain his favorite inference of s_3 using message m_4. If the true state is s_1, any unchecked lie hurts sender 1 (since s_1 is his favorite inference) and so such a lie is not optimal. Finally, suppose the true state is s_2. Obviously, the only lie that could help sender 1 would be to claim s_1 — a claim which must be made using m_1. Notice, however, that sender 2 could refute this claim using m_3 to claim s_3. Since this would lead to the inference s_3, sender 2 would do so and thus sender 1 has no incentive to falsely claim s_1. Obviously, if sender 1 claims the truth, sender 2 cannot refute this claim and so the truth will be inferred. Hence we have a strong separating equilibrium.

Intuitively, this new restriction on the way claims can be made is useful because the key problem for the receiver is to force the senders to reveal that the state is s_2 when this is true. To do so, he requires the "threat" of believing instead that s_3 is true. But if sender 1 sends m_2 in state s_2, the threat to believe s_3 loses its credibility as this report implies that the state cannot possibly be s_3. Hence if the receiver allows this message to be used to claim s_1, sender 1 will falsely claim s_1 when s_2 is true and, given their common interest, sender 2 will not refute this.

This example also illustrates several other interesting points.

1. When full reports exist, requiring the use of full reports may not achieve separation.

Suppose we add a message m_5 which is feasible only in state s_2. This new message technology satisfies (M2) as m_5 is a full report for s_2. Of course, the preference assumption used in Proposition 3 still does not hold. Suppose the receiver uses the BFUR rule — that is, he requires a claim of s to be made using the full report for state s. Surprisingly, there is no separating equilibrium (strong or otherwise) with the BFUR rule. To see why, notice that the full report for state s_1 is m_2. Just as described in the previous paragraph, if a claim of s_1 can be made with this message, the receiver will be misled into inferring s_1 when s_2 is true.

of s_4 is only possible in states s_1 and s_2. As noted, it is never optimal for sender 1 to make a false claim in s_1, so suppose 1 falsely claims s_4 in state s_2. In this case, sender 2 could explicitly challenge to the truth using message m_2. Furthermore, he would certainly do so since this is his favorite inference given that s_3 has been ruled out. A false claim of s_2 is only possible in states s_1 and s_3 (as it must rule out s_4). Again, it could not be optimal in the former state. Obviously, both of the other senders have the ability and incentive to explicitly challenge using message m_3 to claim s_3.

A false claim of s_1 is possible only in states s_2 and s_4. First suppose s_4 is the true state. Sender 2's only options are sending m_1 (*i.e.*, confirming the claim) or explicitly challenging with message m_4 and claiming s_4. (Note that this message *cannot* be used to claim s_2.) Either way, since s_4 is sender 3's favorite inference (given that s_3 has been disproven), he will insure that this is inferred.

The final possibility, then, is that sender 1 claims s_1 in state s_2. This is where the ability to make an implicit challenge is important. Sender 2's options are: m_1 (confirming sender 1's claim), m_4 (explicitly challenging with claim s_4), or m_2 (implicitly challenging with claim s_2). If he uses m_1, sender 3 will use m_4 to explicitly challenge and claim s_4. If he uses m_4, sender 3 will not challenge further. Either way, s_4 is inferred. So suppose he uses the implicit challenge. What are sender 3's options? He cannot refute sender 2's claim since it is true. Hence he can only refute sender 1 (by using m_4) or refute neither claim, thereby generating his least preferred inference, s_2. Either way, then, the final inference is s_2. Since $s_2 \succ_2 s_4$, 2 will use the implicit challenge. Hence we have a strong separating equilibrium.

The key difference between implicit challenges and the kind of inference rules used in the previous results is that sender 1's claim may not be believed even though it is not refuted. More specifically, in the example above, it is important that s_2 be inferred if sender 3 does not refute either claim on the as an *explicit challenge* with this new claim. If he makes no challenge or makes an explicit challenge, then sender 3's options are the same as with the BTUR rule.

Sender 2's new option is what we will call an *implicit challenge*. In an implicit challenge with state s', sender 2 puts s' on the table, ruling out any state s'' such that $M(s'') \subset M(s')$. Unlike an explicit challenge, he does not refute sender 1's claim. Hence following an implicit challenge, there are two states on the table. Implicit challenges are only allowed with states for which an explicit challenge is not possible. More specifically, suppose sender 1 has claimed s_1 using (as required) m_1. Then it is impossible for sender 2 to explicitly challenge with s_2 as he cannot simultaneously refute s_1 and rule out s_4. Hence we allow him to implicitly challenge with s_2 by ruling out s_4 but not refuting sender 1 — that is, he can implicitly challenge with state s_2 using the message m_2. If sender 2 uses an implicit challenge, then at sender 3's turn, there are two states on the table — 1's original claim and 2's implicit challenge. We require sender 3 to refute at least one of these claims. If he does not, the receiver infers whichever of these two claims is worse for 3. If sender 3 refutes one of the claims, he does not make his own claim and the final inference is the unrefuted claim. (With these message sets, 3 could never refute both claims.)

There is a strong separating equilibrium with this inference rule. To see why, first suppose sender 1 claims the true state. Clearly, neither sender 2 nor sender 3 can explicitly challenge the truth. Suppose sender 2 implicitly challenges it. Since sender 3 cannot refute the truth, it is optimal for him to refute 2's claim. Hence the true state will be inferred in this case. Because sender 1 can always generate the correct inference, he always claims s_1 when it is true.

Suppose that sender 1 makes a false claim. Clearly, sender 1 will never falsely claim s_3 since neither of the other senders will refute it (as it is their favorite inference) and this is his least favorite inference. A false claim

2. Banning the use of uninformative messages may preclude separation.

With or without m_5, it is easy to see that the crucial message in achieving separation is m_1. Ironically, this message has *no* pure information content — it proves (rules out) absolutely nothing. However, if it is banned, there is no separating equilibrium. Thus the intuitively plausible conjecture that banning uninformative messages or requiring maximally informative ones helps achieve separation is incorrect. Put differently, the equilibrium information content of apparently uninformative messages can be quite substantial.

3. The existence of a separating equilibrium can depend on the order of the senders.

Consider the original specification of the message technology without m_5. So far, we have only specified preferences over degenerate inferences. Suppose we extend sender 2's preferences to all inferences requiring only that he prefers the degenerate inference s_1 to any other inference. (Since s_1 is his most preferred degenerate inference, this extension seems quite natural.) Then if we reverse the order in which the senders speak, there is no separating equilibrium (strong or not) for any inference rule. To see this, fix any inference rule consistent with feasibility. If there is a separating equilibrium, some combination of messages feasible in state s_1 must lead to the inference s_1. Furthermore, this combination is also feasible in s_2. Let sender 2's message in this combination be m'. So suppose s_2 is the true state. If sender 2 uses the message m', sender 1 will respond with the message that leads to the inference s_1 since this is his favorite possible inference. Hence, since $s_1 \succ_2 s_2$, sender 2 will *not* use any message which leads sender 1 to do something which generates the inference s_2. Thus s_2 cannot be inferred and so there is no separating equilibrium. Intuitively, in state s_2, senders 1 and 2 have a common interest in leading the receiver to falsely infer s_1. When sender 1 goes first, the receiver can exploit the threat of inferring s_3 to deter sender 1 from lying. When sender 2 goes first, sender 1 can always disprove s_3 and so this threat is absent. Hence the common interest prevails.

4. Adding senders can eliminate separating equilibria.

A natural intuition is that the more senders the receiver listens to, the more information he obtains — that is, adding senders creates more informative equilibria. However, this is not true in general. Consider again the original version of the example with sender 1 going first. As noted, there is a strong separating equilibrium. Suppose, though, that we add a third sender with the same preferences as sender 1. Then exactly the argument in point 3 above shows that there is no separating equilibrium. Intuitively, sender 3 can always block the "threat" of inferring s_3 when sender 1 has claimed s_1 in state s_2.

The next example illustrates another approach to modifying the BTUR rule to achieve strong separation with weaker message conditions. In Example 2, we replaced trustworthiness with a stronger criteria. In the following example, we instead introduce a new way to challenge claims.

Example 3.

There are four states, s_1 through s_4, and three senders. The preferences are $s_1 \succ_1 s_2 \succ_1 s_4 \succ_1 s_3$, $s_3 \succ_2 s_2 \succ_2 s_1 \succ_2 s_4$, and $s_3 \succ_3 s_4 \succ_3 s_1 \succ_3 s_2$, so that conflicting preferences holds. The message sets are shown in Figure 3. Notice that (M4) does not hold since $M(s_1) \setminus M(s_3) = \{m_1\} \subset M(s_2)$ and also is a subset of $M(s_4)$. An analogous violation of (M4) arises for s_2.

Consider the following inference rule. Sender 1 can make any trustworthy claim he likes — that is, if he claims s, his message must rule out any states s' such that $M(s') \subset M(s)$. Two types of challenges are available to sender 2. First, as with our earlier inference rules, sender 2 could refute 1's claim, replacing it with any new trustworthy claim. We will refer to this

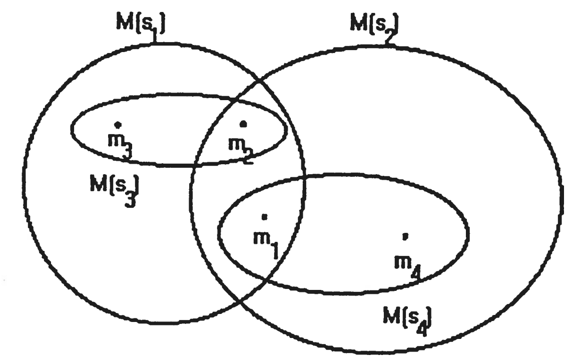

Figure 3.