

Unobserved-Offers Bargaining*

Alexander Wolitzky

MIT

October 3, 2022

Abstract

I study situations where one player (the “claimant”) claims resources, and another player (the “responder”) either accepts or contests the claim. The responder observes the claim only imperfectly. Conflict occurs with positive probability despite perfect information about the players’ preferences. Noisier monitoring of claims predicts greater conflict risk. If the game is repeated, aggressive claims can be deterred and conflict averted if the responder’s signal of the claim is also observed by the claimant. If the responder’s signal is privately observed, it may be impossible to deter the claimant. When both parties make claims, a player is better-off when her own claims are observed more precisely and her opponent’s claims are observed less precisely. Possible applications include international relations, regulation, principal-agency, and product quality provision.

Keywords: bargaining, disagreement, moral hazard, imperfect monitoring, deterrence, salami tactics

JEL codes: C73, C78, D74

*I am grateful for helpful comments from Daron Acemoglu, Sandeep Baliga, Roberto Corrao, Tommaso Denti, Jeffrey Ely, Jim Fearon, Drew Fudenberg, Navin Kartik, Stephen Morris, Doron Ravid, Takuo Sugaya, Jean Tirole, and Muhamet Yildiz, as well as the referees and many seminar participants. I acknowledge financial support from the NSF.

1 Introduction

Most existing theories of disagreement in bargaining are based on imperfect contract enforcement or private information. This paper proposes that the difficulty of making clear offers is another cause of disagreement.

Consider first political and economic conflicts of an asymmetric form, where one party takes an action that establishes some “facts on the ground,” and the other party must then decide whether to acquiesce to this action or to contest it by initiating a costly conflict. International competition often takes this form. China gradually builds military installations on contested islands in the South China Sea; the US decides whether and when to respond. Russia advances in the Caucasus or Ukraine; NATO decides whether to enter the conflict. State-sponsored hackers launch a cyberattack; the victim decides whether to “hack back.”¹ Other examples involve the enforcement of regulatory policies. A factory emits a certain amount of pollution; a regulator decides whether to shut the factory down. A politician posts inflammatory content on social media; the platform decides whether to block her account.² Similar issues also arise in principal-agent relationships. In an efficiency wage contract, a worker chooses an effort level, and her employer decides whether to fire her. In political agency models, a politician chooses a level of misbehavior, such as corruption or election manipulation, and citizens decide whether to oust the politician by voting or protesting.³

These situations have much in common with ultimatum bargaining: one player—the “claimant”—claims a certain amount of resources (physical territory, economic surplus, etc.), and another player—the “responder”—either accepts or rejects the claim, where rejection inefficiently destroys some resources. But they also differ from the standard ultimatum

¹Discussions of these situations in terms of bargaining and deterrence abound. For example, for the South China Sea, see Kaplan (2014) or Coy (2021); for Russia, see Allison (2013) or Freedman (2019); for cyber-deterrence, see Baliga, Bueno de Mesquita, and Wolitzky (2020) and references therein. A less US-centric view of these conflicts would recognize that both sides have opportunities to claim territory and initiate conflicts. In this paper, I consider both asymmetric situations (where one party is the “claimant” and the other is the “responder”) and symmetric ones (where each party plays both roles).

²A large literature in environmental economics studies incentive schemes under imperfect monitoring (Shortle and Horan, 2001). The economics literature on media censorship has so far emphasized a different set of issues (Prat and Strömberg, 2013).

³Unlike in my model, in standard efficiency wage and political agency models (e.g., Shapiro and Stiglitz, 1984; Ferejohn, 1986) the payoff implications of the worker/politician’s past actions are sunk at the time of the firing decision/election, so a wide range of implicit contracts are credible for the employer/citizens.

bargaining game in two key ways. First, the responder observes the claim only imperfectly: that is, he does not know exactly what payoff he will receive if he accepts the claim. (Is an installation designed for civilian or military use? Is an inflammatory Tweet hyperbole or a genuine incitement to violence? How much pollution will a factory emit?) I model this feature by assuming that the responder observes an informative signal of the claim, which satisfies a full-support assumption. Second, the claimant can make a claim that is so aggressive that the responder strictly prefers to reject it: formally, the claimant can demand more than the whole pie. This possibility is ignored in standard bargaining models without loss of generality, because such offers would always be rejected when offers are perfectly observed; however, it must be taken into account when offers are imperfectly observed.⁴

My first set of results characterizes the Nash equilibria of this game. Since the responder’s signal has full support, there are always trivial equilibria where the claimant demands more than the whole pie and the responder rejects following any signal realization.⁵ I call an equilibrium *non-trivial* if the responder accepts with positive probability. Non-trivial equilibria exist if the responder’s signal is sufficiently informative. In every non-trivial equilibrium, the claimant demands the whole pie (but no more), leaving the responder indifferent between accepting her demand and rejecting it. In the Pareto-optimal non-trivial equilibrium (i.e., the equilibrium with the highest probability of acceptance), the responder accepts if and only if the signal s falls below a cutoff s^* . Thus, as in standard ultimatum bargaining, the claimant demands the entire surplus; but unlike the usual case, this demand is rejected with positive probability. Intuitively, the responder cannot always accept in equilibrium, or else the claimant would increase her demand beyond the responder’s reservation utility, which would cause the responder to reject. The model thus predicts a positive probability of inefficient conflict, even though the bargaining surplus is perfectly divisible and the players’ preferences are common knowledge.⁶ The model is also tractable and has intuitive com-

⁴Ravid (2020) and Denti, Marinacci, and Rustichini (2022) study ultimatum bargaining with a rationally-inattentive responder and with costly information acquisition, respectively; I discuss their papers in Section 2. A different role for unobserved offers arises in multilateral bargaining with externalities, where bargainers must form beliefs about the agreements reached by third parties (e.g., Rey and Tirole, 2007).

⁵This logic is as in Bagwell (1995).

⁶From the perspective of international relations theory, none of Fearon’s (1995) “rationalist explanations for war” seem to apply in my model. So “imperfectly observed offers” may be another such explanation. In particular, my model does not involve “private information” in the sense of a payoff-relevant move by

parative statics. For example, the probability of conflict decreases as the responder’s signal becomes more precise, and it vanishes in the perfect-observability limit, where the Pareto-efficient equilibrium converges to the unique subgame-perfect equilibrium of the standard ultimatum bargaining game.

The model expresses a view of inefficient conflict based on imperfect observability and indifference. Since the claimant claims the whole pie, the responder is willing to start a conflict. If the claimant could claim a penny less—and could perfectly communicate this claim to the responder—conflict would be averted, leaving both parties better-off. But the claimant cannot perfectly communicate her claim, because the claim itself is not perfectly observed; moreover, allowing the claimant to send cheap talk messages to the responder does not help.⁷ Given that the claim cannot be perfectly communicated, it is not optimal for the claimant to demand a penny less, because the resulting small reduction in conflict risk is worth less than a penny. Finally, not only is the responder willing to reject the claim in every equilibrium, but rejection must actually occur with positive probability in every equilibrium: indeed, I focus on the equilibrium with the smallest rejection probability.

With the analysis of this simple, one-shot bargaining game in hand, I next consider the richer model where the bargaining game is played repeatedly. The repeated model is more realistic—since in reality parties can change their claims over time—and it also introduces the question of whether repeated game effects can deter aggressive claims and reduce the risk of conflict. I show that the answer to this question depends on whether the responder’s signal s is also observed by the claimant: that is, on whether s is *public* or *private*. With sufficiently informative public signals, a folk theorem can be supported by trigger strategies of a very simple form: the claimant claims a share $x < 1$ of the pie; the responder accepts if and only if s falls below a cutoff s^* ; and if the responder deviates by accepting when $s > s^*$ (which is profitable, given that $x < 1$), this triggers a switch to permanent conflict with positive probability. This equilibrium resembles the “trip wire” or “plate glass window” strategies that Schelling (1966, p. 47) suggested were essential for deterrence: the cutoff

Nature—it is a moral hazard model, not an adverse selection model. Note that, as in standard moral hazard models (e.g., Holmström, 1979), the claim/action is “known” in equilibrium despite being imperfectly observed.

⁷Intuitively, cheap talk does not help because statements that the claim is low are not credible. More credible methods of communication could be valuable for the claimant—this is left for future research.

signal s^* is a “line in the sand” that must spark conflict when it is crossed, or else both parties realize that deterrence has broken down. In contrast, with private signals, deterrence is impossible for a large class of strategies, including any strategy profile where the claimant’s strategy depends only on the history of the responder’s decisions and the responder’s strategy depends only on this history in addition to the current-period signal. However, while the claimant cannot be deterred from claiming the whole surplus, the probability of rejection can be made arbitrarily small when the players are patient. Intuitively, with private signals, if $x < 1$ then the responder accepts after every signal, which is inconsistent with equilibrium; but if $x = 1$ —so that any cutoff signal is consistent with optimal play by the responder—then when the claimant is patient there is a cutoff signal that deters her from claiming $x > 1$ but also triggers rejection with only a small probability.

The repeated bargaining model gives the claimant complete flexibility to change her demands over time, and similarly lets the responder switch back and forth between accepting and rejecting these demands. In some settings, it seems more realistic to instead assume that the claimant’s demand can only increase over time, or to assume that rejecting a claim leads to a permanent state of conflict. I thus also consider variants of the model with monotone claims or irreversible rejection. These variants may be viewed as models of “salami tactics” (Schelling, 1966), where one party makes increasingly aggressive demands, and its adversary decides whether and when to fight.⁸ The main predictions about deterrence and conflict risk in these model variants turn out to be exactly the same as in the repeated bargaining model. However, the irreversible-rejection model yields a novel, pessimistic prediction regarding the long-run outcome of the game: the probability that the players find themselves in conflict in period t —that is, the probability that a claim was rejected in some period prior to t —converges to 1 as $t \rightarrow \infty$.

Finally, to capture more symmetric settings where both parties have opportunities to claim resources, I consider a version of the one-shot bargaining model where the parties take turns playing the roles of claimant and responder, as in Rubinstein’s (1982) alternating-offers bargaining model. My analysis of unobserved-offer ultimatum bargaining extends naturally

⁸Powell (1996a) models salami tactics as an incomplete-information war of attrition; my complete-information bargaining model captures a different aspect of salami tactics. The assumption that rejection ends the game is as in “crisis-bargaining” models like that of Fearon (1994) or Powell (1996b).

to this alternating unobserved-offers game. A new result is that, for a natural equilibrium selection that generalizes the Pareto-optimal equilibrium of the baseline model (“extremal stationary cutoff equilibria”), a player’s payoff is higher when her opponent observes her demand more precisely, as well as when she observes her opponent’s demand less precisely. The ability to make clear demands—and to avoid perceiving the opponent’s demand clearly—is thus a strategic advantage in bargaining.⁹ For example, a firm with skilled lawyers and a busy CEO may benefit from offering its partners clear contracts, while enjoying a credible commitment not to pay much attention to any counteroffers.

The paper is organized as follows: Section 2 introduces and analyzes the one-shot bargaining game. Section 3 studies repeated bargaining, comparing the cases with public and private signals, and also considers extensions to monotone claims and irreversible rejection. Section 4 analyzes alternating unobserved offers. Section 5 concludes. Proofs are deferred to the appendix.

2 Ultimatum Bargaining with Imperfect Observation

2.1 The One-Shot Bargaining Game

A *claimant* (“she”) and a *responder* (“he”) bargain over a unit surplus. The timing is as follows: First, the claimant chooses a *claim* (or *demand*) $x \in X = [-M, M]$, where M is a parameter strictly greater than 1. The responder then observes a *signal* $s \in S \subseteq \mathbb{R}$ of x , where the distribution of s conditional on x is denoted $F(s|x)$. Finally, the responder decides whether to accept or reject the claim. If the responder accepts, payoffs are x for the claimant and $1 - x$ for the responder. If the responder rejects, each player’s payoff is 0.

I assume that the signal space S is a closed interval, which can be bounded or unbounded, so $S = [\underline{s}, \bar{s}]$ with $\underline{s} \in \mathbb{R} \cup \{-\infty\}$ and $\bar{s} \in \mathbb{R} \cup \{+\infty\}$. I impose the following standard assumptions on F .

Assumption 1 The signal distribution F satisfies

⁹Schelling (1960, p. 26) also noted a version of this point: “An asymmetry in communication may well favor the one who is (and is known to be) unavailable for the receipt of messages, for he is the one who cannot be deterred from his own commitment by the receipt of the other’s.”

1. *Smoothness*: F admits a density, f , which is twice continuously differentiable in x .
2. *Full support*: $f(s|x) > 0$ for all s, x .
3. *Strict Monotone Likelihood Ratio Property (MLRP)*: $f(s|x)f(s'|x') > f(s|x')f(s'|x)$ for all $s < s', x < x'$.
4. *Log-concavity in x* : $F_x(s|x)/F(s|x) \geq F_x(s|x')/F(s|x')$ for all $s, x < x'$.

Full support is the key difference from standard ultimatum bargaining. MLRP says that higher signals are informative of higher claims, and are thus “bad news” for the responder (Milgrom, 1981). Log-concavity will imply that the claimant’s problem is quasi-concave; it is satisfied by many common probability distributions (Bagnoli and Bergstrom, 2005).

A strategy for the claimant is a probability distribution over X . A strategy for the responder specifies, for each signal $s \in S$, a probability $\sigma(s)$ of accepting. Since F has full support, every signal is on-path, so ex ante and sequential rationality coincide for the responder. I thus use the Nash equilibrium (NE) solution concept. Also, since F admits a density, I identify responder strategies that differ only on a measure-0 set of signals.

I call an equilibrium *trivial* if the responder rejects with probability 1. I will show that the model has a non-trivial equilibrium if and only if the following condition holds.

Informativeness Condition $\lim_{s \rightarrow \underline{s}} f_x(s|x=1)/f(s|x=1) < -1$.

Intuitively, this condition says that when the demand is concentrated around $x = 1$, a sufficiently low signal is sufficiently good news about the demand.

Example 1 *All of the above assumptions, including the informativeness condition, are satisfied if $s = x + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \theta^2)$ is a mean-zero normal random variable, with variance θ^2 fixed independently of x .*

A few remarks are in order. First, the assumption that $M > 1$ says that the claimant can demand more than the whole pie. This is a crucial assumption. However, so long as it is greater than 1, the precise value of M is irrelevant. Second, for the current, one-shot model, it makes no difference whether or not the claimant observes s . However, this distinction will

matter when I consider repeated bargaining in Section 3. Third, the terminology “one-shot bargaining” serves to contrast the current model (which I also call the “baseline model”) with the repeated bargaining model of Section 3 and the alternating-offers bargaining model of Section 4. Of course, the one-shot bargaining game is itself a dynamic game.

2.2 Equilibrium

My first result characterizes equilibria in the one-shot bargaining game.

Theorem 1 *A strategy profile is a Nash equilibrium of the one-shot bargaining game if and only if it takes one of the following two forms:*

1. (Trivial) *The claimant’s strategy satisfies $\mathbb{E}[x|s] \geq 1$ for all $s \in S$. The responder always rejects.*
2. (Non-Trivial) *The claimant always demands $x = 1$. The responder’s strategy σ satisfies $1 \in \operatorname{argmax}_{x \in X} x \int_{s \in S} \sigma(s) f(s|x) ds$.*

Moreover, when a non-trivial equilibrium exists, there is a unique non-trivial equilibrium that maximizes the probability that the responder accepts (and thus is Pareto optimal): in this equilibrium, the responder accepts if and only if $s \leq s^$, where $s^* > \underline{s}$ satisfies*

$$F(s^*|x=1) + F_x(s^*|x=1) = 0. \tag{1}$$

Finally, a non-trivial equilibrium exists if and only if the informativeness condition holds.

In a trivial equilibrium, the claimant “usually” makes unacceptable demands, so the responder rejects after every signal. For example, for any $x \geq 1$, it is an equilibrium for the claimant to always demand x while the responder always rejects. In a non-trivial equilibrium, the claimant demands $x = 1$ with probability 1, so the responder gets zero surplus and is indifferent between accepting and rejecting after every signal. However, not every responder strategy is consistent with equilibrium: in order for the claimant’s demand of $x = 1$ to be optimal, the responder must use a strategy of the form described in the theorem. The non-trivial equilibria are clearly the interesting ones (when the signal is informative enough that

they exist). In the next subsection, I will also argue that there are good theoretical reasons to focus on the Pareto-optimal non-trivial equilibrium characterized in the theorem, where the responder uses a cutoff strategy.¹⁰

The most subtle points in Theorem 1 are that there cannot be a non-trivial mixed-strategy equilibrium (and hence claims other than $x = 1$ are never accepted), that the Pareto-optimal equilibrium takes a cutoff form, and that a non-trivial equilibrium exists iff the informativeness condition holds. There is no non-trivial mixed-strategy equilibrium because if the claimant mixes then $\mathbb{E}[x|s]$ is strictly increasing in s (as a consequence of MLRP), so the responder must use a cutoff strategy; however, the claimant has a unique best response to any cutoff strategy with cutoff $s^* \in (\underline{s}, \bar{s})$ (as a consequence of log-concavity), and hence cannot mix. The Pareto-optimal equilibrium takes a cutoff form because maximizing the responder’s acceptance probability subject to the requirement that $x = 1$ is an optimal demand for the claimant entails setting $\sigma(s) = 1$ if $f_x(s|x)/f(s|x) < \lambda$ and $\sigma(s) = 0$ if $f_x(s|x)/f(s|x) > \lambda$ for some threshold λ , which yields a cutoff form under MLRP. Finally, a non-trivial equilibrium exists iff there is a cutoff equilibrium with $s^* > \underline{s}$. Under MLRP, this is the case iff the informativeness condition holds.

The Pareto-optimal equilibrium in my model is related to an equilibrium in Ravid (2020). Ravid studies ultimatum bargaining with a rationally-inattentive responder (buyer), focusing on the case where the claimant (seller) is privately informed, so the responder pays an “attention cost” to simultaneously learn about the claimant’s demand and her type. When the claimant is not privately informed (Ravid’s “known-quality case,” where the responder only needs to learn about the demand), in the unique equilibrium satisfying a trembling-hand-like refinement, optimal attention implies a “recommendation strategy” for the responder with the same behavioral implications as observing $s = x + \varepsilon$, with ε logistically distributed,

¹⁰The pure-strategy trivial equilibria of the one-shot bargaining game (i.e., the trivial equilibria where the claimant always demands some $x \geq 1$) are precisely the pure-strategy NE of the simultaneous-move game where the claimant chooses $x \in X$ and the responder chooses accept or reject. These strategy profiles remain NE in the dynamic game because the signal has full support, by the argument of Bagwell (1995). The dynamic game can also have additional pure-strategy NE—the pure-strategy non-trivial equilibria—because the responder has multiple best responses to $x = 1$. This contrasts with Bagwell’s analysis, which assumes a unique best response for the second-mover. Of course, multiple best responses are endemic in bargaining games. For example, in the unique subgame-perfect equilibrium of the standard ultimatum game, the proposer demands $x = 1$, so both accept and reject are best responses.

and using a cutoff strategy. Since this signal distribution satisfies my assumptions, the same argument as in the proof of Theorem 1 then implies that the claimant always demands $x = 1$.

Denti, Marinacci, and Rustichini (2022) compare rational inattention and costly information acquisition in general decision problems. As an application, they consider ultimatum bargaining with an unobserved demand (which may exceed the size of the pie), where the responder must pay a strictly positive cost to get any information about the demand at all. This “no free information” assumption is violated in my exogenous-signal model. With no free information, disagreement must arise with positive probability (as in my model), and the responder obtains a positive surplus in any non-trivial equilibrium (unlike in my model, where the responder gets zero surplus). Moreover, the same argument as in the proof of Theorem 1 shows that *only* trivial equilibria can exist under no free information if all acquirable information structures satisfy MLRP and log-concavity, since these assumptions preclude the existence of a mixed equilibrium. Denti, Marinacci, and Rustichini construct a non-trivial mixed equilibrium for a particular class of information structure. By Theorem 1, this class must violate one of my assumptions; indeed, it violates log-concavity.

Several authors have studied whether introducing communication reduces the probability of disagreement in various bargaining models (e.g., Farrell and Gibbons, 1989; Fearon, 1995). It is therefore worth noting that Theorem 1 holds verbatim when the game is augmented by letting the claimant make a cheap talk statement after choosing her demand.¹¹ The reason why is that for every message m that the claimant may send which leaves some uncertainty about x , the responder will accept iff $s \leq s^*(m)$ for some message-contingent cutoff $s^*(m)$. The claimant will therefore only send messages that induce the highest cutoff, so the message cannot convey any useful information about the demand.

Proposition 1 *Theorem 1 remains valid when the game is augmented with cheap talk.*

It is also natural to ask what happens if the claimant can offer a perfectly observable cash payment of y dollars in addition to the unobserved claim x , so that payoffs on acceptance are $x - y$ for the claimant and $1 - x + y$ for the responder. This extension captures situations where the claimant’s proposal has multiple components, some of which are perfectly observed: for

¹¹Here $\sigma(s)$ must be viewed as the probability that the responder accepts conditional on signal realization s , integrating over the claimant’s message.

example, a cash transfer or the location of a physical boundary may be perfectly observable, while compliance with an arms control or environmental agreement is not. Suppose that the claimant can offer any cash payment up to some maximum amount L . If $L < M - 1$, one can simply define the “all-in” claim to be $\tilde{x} = x - y$, and the above analysis goes through verbatim, because the all-in claim can strictly exceed 1 even if y takes its maximum value of L . If instead $L \geq M - 1$, then it is an equilibrium for the claimant to set $(x = M, y = M - 1)$ while the responder accepts after every signal. Thus, once one allows a perfectly observed component of the claimant’s proposal, my analysis applies if and only if the range of perfectly observable concessions that the claimant can make is smaller than the range of imperfectly observable claims.

2.3 Selecting the Pareto-Optimal Equilibrium

Theorem 1 implies that the one-shot bargaining game can have many equilibria. However, the trivial equilibria are uninteresting, and the non-trivial equilibria where the responder’s strategy does not take a cutoff form seem rather artificial. Beyond these impressionistic reasons for focusing on the Pareto-optimal equilibrium (where the responder accepts iff $s \leq s^*$ and s^* satisfies equation (1)), this subsection gives two theoretical arguments for doing so.

First, note that if $F(s|x) = \mathbf{1}\{s \geq x\}$ —so the claim is perfectly observed—the one-shot bargaining game reduces to the standard ultimatum game. In the unique subgame-perfect equilibrium in this limit game (also called the *Stackelberg outcome*), the claimant demands $x = 1$ and the responder accepts iff $s \leq 1$. Thus, an equilibrium outcome of the one-shot bargaining game is continuous in the limit as observation noise vanishes iff it converges to the Stackelberg outcome. Van Damme and Hurkens (1997) call such outcomes *noisy Stackelberg equilibria*. If an analyst predicts the subgame-perfect/Stackelberg equilibrium in the limit game and wants her predictions to be robust to introducing a small amount of observation noise, she must predict a noisy Stackelberg equilibrium in the imperfect monitoring game.

I show that as observation noise vanishes, the Pareto-optimal equilibrium of the one-shot bargaining game converges to the subgame-perfect equilibrium of the ultimatum game. That is, the Pareto-optimal equilibrium is the noisy Stackelberg equilibrium.

Proposition 2 *Suppose that $S \supseteq X$. For any sequence of signal distributions $\{F^n(s|x)\}_{n \in \mathbb{N}}$ satisfying $\sup_{s,x} |F^n(s|x) - \mathbf{1}\{s \geq x\}| \rightarrow 0$ and $f_x^n(0|x=1)/f^n(0|x=1) < -1$ for all n , and any corresponding sequence of equilibrium cutoff signals $\{s^{n,*}\}_{n \in \mathbb{N}}$ satisfying $s^{n,*} > \underline{s}$ for all n , we have $s^{n,*} \rightarrow 1$ and $F^n(s^{n,*}|1) \rightarrow 1$.*

For example, the sequence of signal distributions may be given by $s = x + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \theta^2)$ and $\theta^2 \rightarrow 0$. Note that since $S \supseteq X$, the condition $f_x(0|x=1)/f(0|x=1) < -1$ is a strengthening of the informativeness condition.

A second argument is that in many applications it is reasonable to think that the signal itself is slightly payoff-relevant, so that the responder's payoff from accepting the claim is not $1 - x$ but $1 - (\pi x + (1 - \pi)s)$, where $\pi \in (0, 1)$ is a new parameter. For example, suppose that x is the amount by which a new factory deviates from an environmental regulatory standard, and s is amount of pollution emitted by the factory in its first month of operation. Since the factory's emissions over the long run are determined by x rather than s , in deciding whether to shut down the factory a regulator mostly cares about x , and is thus interested in s primarily through its information content concerning x . But perhaps s rather than x determines the factory's likely emissions in its second month of operation, so s is also slightly payoff-relevant for the regulator. I refer to situations where the responder's payoff from accepting is $1 - (\pi x + (1 - \pi)s)$ as having *payoff-relevant signals*.¹²

Proposition 3 *Assume that the informativeness condition holds. With payoff-relevant signals and $S = \mathbb{R}$, there is a unique NE, and it converges to the Pareto-optimal equilibrium with payoff-irrelevant signals as $\pi \rightarrow 1$.*

The intuition is that the responder must use a cutoff strategy when $\pi < 1$, and as $\pi \rightarrow 1$ this strategy converges to the (unique) equilibrium cutoff strategy in the baseline model.

Making signals slightly payoff-relevant is one way of perturbing the model to make the responder's best response strict. It should be noted that other perturbations select different

¹²Here I assume that the claimant's payoff when the responder accepts remains x , rather than $\pi x + (1 - \pi)s$. The same result applies for the latter specification, so long as F is such that $\mathbb{E}[(\pi x + (1 - \pi)s) \mathbf{1}\{\pi x + (1 - \pi)s \leq s^*\}]$ has a unique maximizer for any s^* . As well as being slightly simpler, the former specification is more natural in applications like the pollution example, where the factory owner benefits from skirting environmental regulations but does not directly benefit from emitting pollution. The general approach of perturbing the game by making s payoff relevant is somewhat in the spirit of Maggi (1999), who studied a variant of Bagwell (1995) where the first-mover has private information.

equilibria. For example, if the responder’s payoff from accepting the claim is perturbed to $1 - x + \rho$, where ρ is an independent “taste shock,” then the model has a unique pure-strategy equilibrium, where the claimant always demands $x = M$ and the responder accepts iff $\rho \geq M - 1$. The conclusion of this subsection is thus that some attractive approaches to equilibrium selection pick out the Pareto-optimal equilibrium, not that every plausible approach does so.¹³

2.4 Varying the Amount of Observation Noise

Recall that the disagreement probability converges to 0 as observation noise vanishes (Proposition 2). A natural comparative static is that this convergence is monotone if noise enters the signal additively and is reduced by a scalar factor.

Proposition 4 *Suppose that $s = x + \theta\varepsilon$, where $\theta > 0$ is a parameter and ε is a random variable whose distribution G admits a density g satisfying Assumption 1. Let s^* be given by equation (1). Then $F(s^*|x = 1)$ is decreasing in θ .*

Since the claimant’s equilibrium expected utility equals $F(s^*|x = 1)$, Proposition 4 implies that the claimant always benefits from making her demand more transparent, while the responder’s equilibrium expected utility is 0 regardless of the transparency of the claimant’s demand. In Section 4, I show that with alternating-offers bargaining, each party’s expected utility is greater when their own demand is more transparent, as well as when their opponent’s demand is less transparent.

A slightly subtle aspect of Proposition 4 arises when $s^* < 1$. In this case, the “direct effect” of increasing θ (i.e., the effect holding s^* fixed) is to *increase* the acceptance probability. Proposition 4 says that, when $s^* < 1$, this direct effect is always outweighed by an indirect effect on the acceptance probability that enters through a decrease in s^* . The intuition is that the claimant’s marginal benefit of increasing x equals the acceptance probability

¹³An interesting open question is whether the additive taste-shock model has a mixed equilibrium that converges to the Stackelberg outcome as observation noise vanishes. Note that such an equilibrium *cannot* exist when $xF(s|x)$ is concave in x for each s , as in this case the claimant’s best-response is unique for any distribution over cutoffs s^* .




Figure 1: Agreement Probability as a Function of Observation Noise

$F(s^*|x=1)$, while her marginal cost of increasing x equals $-F_x(s^*|x=1)$; and since the marginal cost is decreasing in θ , the marginal benefit must be so as well.

Figure 1 illustrates Proposition 4 in the normal-noise case, $G = \mathcal{N}(0, \theta^2)$. The figure graphs the agreement probability as a function of θ^2 : the dependent variable is the value of $\Phi_{\theta^2}(s^*)$ at which $\Phi_{\theta^2}(s^*) = \phi_{\theta^2}(s^*)$, where ϕ_{θ^2} and Φ_{θ^2} are the pdf and cdf of a $\mathcal{N}(0, \theta^2)$ random variable. For example, the agreement probability is 70% if $\theta^2 = 0.5$, but falls to 38% when $\theta^2 = 1$.

2.5 Discussion

I briefly discuss how my model relates to classic bargaining models with incomplete information or “noisy” demands, as well as possible dynamic interpretations of the model.

A standard explanation of disagreement in bargaining is that the proposer/claimant does not know the responder’s reservation value. Unknown reservation values and unobserved offers both lead to disagreement with positive probability, and in both models the probability of disagreement is increasing in the amount of uncertainty (under appropriate conditions).¹⁴ Other predictions differ between the models. With unobserved offers the claimant’s expected

¹⁴If the responder’s reservation value is $v = v_0 + \theta\varepsilon$, where $\theta > 0$ is a parameter and ε is a random variable, then the probability of agreement is decreasing in θ if $1 - F(v)$ is log-concave, similarly to Proposition 4.

payoff is always decreasing in the amount of observation noise, while with unknown reservation values the effect of uncertainty on the claimant’s expected payoff is ambiguous.¹⁵ Also, with unobserved offers the claimant demands 1 and the responder’s expected payoff is 0, regardless of the distribution of observation noise, while with unknown reservation values the claimant’s demand and the responder’s (positive) expected payoff depend on the distribution of reservation values. The distinctive feature that with unobserved offers the responder is always “on the fence” between accepting and rejecting fits many of my motivating examples.

Other differences between the two models concerns dynamic bargaining. If the proposer can quickly change her demand when it is rejected, the Coase conjecture (Fudenberg, Levine, and Tirole, 1985; Gul, Sonnenschein, and Wilson, 1986) predicts that disagreement in incomplete-information bargaining is short-lived, and a similar conclusion holds in repeated incomplete-information bargaining with a persistent reservation value (Hart and Tirole, 1988). It is thus unclear if incomplete information about reservation values is a robust explanation for persistent disagreement. In contrast, it seems natural to assume that observation noise is iid across periods in the repeated version of my model, in which case the only Markov equilibria are repetitions of one-shot equilibria, which feature a constant probability of disagreement.¹⁶ Also, alternating-offers bargaining with two-sided uncertainty about reservation values is notoriously difficult to analyze (Ausubel, Cramton, and De-neckere, 2002), while I show in Section 4 that the alternating-offers version of my model is tractable and yields some new insights.

Another relevant comparison is to Nash’s demand game (Nash, 1953; Binmore, 1987; Abreu and Pearce, 2015), where random perturbations of the bargaining set are used to select the Nash bargaining solution. Carlsson (1991) analyzes a similar model, where the players’ demands are perturbed rather than the bargaining set. In Carlsson’s model, a demand of 50 cents may be perturbed to 51 cents, while in my model a demand of 50 cents may be misperceived as 51 cents, but is always “really” (in terms of the final payoffs on acceptance) 50 cents. The apparently minor difference of whether noise hits a player’s

¹⁵If $v = v_0 + \theta\varepsilon$, where $\theta > 0$ is a parameter and ε is a mean-0 random variable, the effect of an increase in θ on the proposer’s expected utility can be positive or negative, depending on the distribution of ε , in contrast to Proposition 4. To see why, note that changing θ rotates the “demand curve” facing the proposer around v_0 , which can either increase or decrease profit (cf. Johnson and Myatt, 2006).

¹⁶Non-Markov equilibria in this model are the subject of Section 3.

demand itself or the opponent's signal of the demand is actually important, because the opponent's interpretation of the signal is endogenously determined by Bayes' rule. Another difference is that Nash's game involves simultaneous demands. For these reasons, the analysis and conclusions of Nash's demand game and its variants are quite different from mine.

Finally, I note that my baseline model can be given a dynamic interpretation, where time runs from $t = 0$ to M , and the claimant gradually claims resources over time at rate 1, stopping when she has claimed x resources. At time $t = M$, the responder gets a signal s of the claim x , and chooses accept or reject. Payoffs are realized at the end of the game. This dynamic interpretation is similar to the truly dynamic model with monotone demands and irreversible rejection considered in Section 3.3, with the differences that in the genuinely dynamic model the claimant claims resources at an endogenous and possibly history-dependent rate, the responder can reject at any time (ending the game), signals and payoffs realize over time rather than only at the end of the game, and the claimant may (depending on the model specification) observe the signals. As we will see, these differences matter: they lead the dynamic model to have a larger set of equilibria than the baseline model, possibly including equilibria where the responder obtains a positive payoff. Nonetheless, the baseline model already captures some situations where claims accrue over time.

2.6 Extensions

The baseline model is so stylized that its fit with any particular application is somewhat loose. To conclude this section, I sketch two extensions that fit economic applications more tightly. The goal here is illustrating the model's applicability rather than studying the applications per se, so I only sketch the analysis. Nonetheless, the model does seem to deliver some insights in these settings.

2.6.1 Uncertain Contract Interpretation

Suppose that the claim is a formal contract offered to the responder. If the responder accepts, the contract will eventually be presented to a court, which interprets it. The court's interpretation is uncertain at the time of contracting. The parties' final payoffs are determined

by the court’s interpretation of the contract.

To model this situation, let x equal the expected interpretation of the contract, from the perspective of the claimant at the moment she writes the contract. The court’s interpretation is given by $x + \varepsilon_1$, where ε_1 is a mean-zero random variable, and the final payoffs if the contract is accepted are given by $x + \varepsilon_1$ for the claimant and $1 - x - \varepsilon_1$ for the responder. When deciding whether to accept the contract, the responder observes $x + \varepsilon_2$, where ε_2 is another random variable. This corresponds to the responder’s own interpretation of the contract, which may differ from the court’s interpretation.¹⁷

This model with a general correlation structure between ε_1 and ε_2 —the court’s interpretation and the responder’s interpretation—can be analyzed along the lines of the baseline model. Two special cases bear emphasis. First, when ε_1 and ε_2 are independent, we recover the baseline model, as in this case ε_1 is just an independent mean-zero shock to the players’ payoffs when the contract is accepted, which does not affect incentives since the parties are risk-neutral. Second, when ε_1 and ε_2 are perfectly correlated, we recover a standard lemons model, where $\varepsilon_2 (= \varepsilon_1)$ is the responder’s payoff-relevant private information.¹⁸

The uncertain-contract-interpretation model thus interpolates between the baseline model and a lemons model, where the baseline model arises when the responder has no private information about the court’s interpretation, and the lemons model arises when the responder can perfectly predict it. One can ask various questions of this model, for example deriving conditions under which lower variation in ε_2 leads to a higher probability of agreement, as is true in the polar cases of the baseline model and (under similar conditions) the lemons model.

2.6.2 Moral Hazard in Product Quality

Suppose that the claimant is a seller who produces a good of quality q at cost $c(q)$, and offers it for sale to the responder (buyer) at a price p . The buyer observes p perfectly, but

¹⁷Uncertain contract interpretation is somewhat related to costly contemplation as modeled by Tirole (2009), Bolton and Faure-Grimaud (2010), and Ravid (2020).

¹⁸To spell out the mapping, suppose that the claimant is a buyer who offers price p for a good of uncertain quality q , where q is the responder’s (seller’s) private information, and a quality- q good is worth $1 + q$ to the buyer and q to the seller. Letting $x = 1 - p$ and $\varepsilon_2 = q$ recovers the current model, as payoffs on acceptance are $1 + q - p = x + \varepsilon_2$ for the buyer and $-q + p = 1 - x - \varepsilon_2$ for the seller.

observes only a noisy signal $s \sim F(s|q)$ of q . If the buyer accepts, payoffs are $p - c(q)$ for the seller and $q - p$ for the buyer; if the buyer rejects, payoffs are $-c(q)$ and 0, respectively. Thus, the production cost $c(q)$ is sunk at the time of contracting, and the buyer's valuation of a quality- q good equals q .¹⁹

Let us impose a refinement that equilibria of this game, where the seller chooses q and p simultaneously, should also be equilibria of the game where the seller chooses first p and then q . (This rules out equilibria where the seller is forced to name a particular price by threatening that different prices induce adverse beliefs about quality.) Then, since the price p is observed, the model is equivalent to the baseline model with a “perceived offer” to the buyer of $\mathbb{E}[q|s] - p$, but where the seller must incur a sunk cost to improve the perceived offer. By the same logic as in the baseline model, there is a unique non-trivial pure-strategy equilibrium (which is also the unique non-trivial equilibrium overall under appropriate conditions, given the aforementioned refinement). In this equilibrium, the seller sets $q = p$, and the buyer is indifferent but nonetheless accepts or rejects based on s . Making quality more observable increases the probability of acceptance. Interestingly, quality can be above or below the first-best level (given by $c'(q) = 1$). The intuition for why quality can be above first-best is that naming a high price has a commitment benefit for the seller, as it improves her incentive to make a sale and thus induces her to provide high quality. In this case, counterintuitively, making quality more observable decreases quality and price (toward the first-best level).

3 Repeated Bargaining: Prospects for Deterrence

Now suppose the one-shot bargaining game is played repeatedly over an infinite horizon. The players maximize their discounted expected payoffs, with discount factor $\delta \in [0, 1)$. For convenience, I also assume that the players observe the outcome of a public randomizing device $z_t \sim U[0, 1]$ at the end of every period t . I consider two versions of this repeated-game

¹⁹There is a vast literature on moral hazard in product quality, including major branches on reputation (e.g., Klein and Leffler, 1981; Milgrom and Roberts, 1986) and search (e.g., Wolinsky, 1983). However, I have been surprisingly unable to locate work where quality-provision is modeled as a one-shot moral hazard problem with imperfect monitoring, as in the current setting.

model: *public signals*, where the period- t signal $s_t \sim F(\cdot|x_t)$ is observed by the claimant as well as the responder; and *private signals*, where s_t is observed only by the responder. With public signals, the *public history* (i.e., the history of all publicly available information) at the beginning of period t is $h^t = (s_\tau, a_\tau, z_\tau)_{\tau=1}^{t-1}$, where $a_\tau \in \{\text{accept, reject}\}$ is the responder's period- τ decision, while with private signals $h^t = (a_\tau, z_\tau)_{\tau=1}^{t-1}$. I briefly discuss richer information structures where the players observe correlated signals in Section 3.2.²⁰

A player's behavioral strategy is *public* if it depends only on the public history h^t . That is, a public strategy for the claimant specifies a distribution over X for each h^t , and a public strategy for the responder specifies an acceptance probability for each pair (h^t, s_t) , where $s_t \in S$ is the current-period signal. I focus on the solution concept of *public perfect equilibria* (*PPE*), which is a profile of public strategies that form a Nash equilibrium conditional on any public history.²¹ Note that this solution concept is well-defined with both public signals and private signals—the difference between the two versions of the model is whether the past signals $(s_\tau)_{\tau=1}^{t-1}$ are included in h^t , not the solution concept. PPE is a standard solution concept in repeated games, but it is not without loss of generality. In what follows, I discuss how my results might change under a more permissive solution concept.

3.1 Public Signals

I show that a simple type of trigger strategies can be used to deter aggressive claims and avert conflict, so long as the signal s is publicly observed and satisfies the following condition:

Unbounded Informativeness For all $x \in X$, $\lim_{s \rightarrow \bar{s}} f_x(s|x) / f(s|x) = \infty$.

Intuitively, this condition says that sufficiently high signals are unboundedly strong bad news about the demand. Again, this condition holds if $s = x + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \theta^2)$.

Definition 1 A trigger strategy profile is characterized by numbers $x \in (0, 1]$, $s \in S$, and $\alpha \in [0, 1]$, and is defined as follows:

²⁰In this section, the responder's payoff is not measurable with respect to his information. This could be remedied by introducing a second signal of x_t , which is observed at the end of period t . Adding such a signal does not affect the results of this section, because the strategies I consider would not condition on it.

²¹PPE were introduced by Fudenberg, Levine, and Maskin (1994). Of course, the Nash equilibrium requirement allows players to contemplate deviations to any strategy, including non-public ones.

1. In Phase 1, the claimant demands x and the responder accepts iff $s_t \leq s$.
2. In Phase 2, the claimant demands 1 and the responder rejects after every signal.
3. Play begins in Phase 1, and permanently transitions to Phase 2 at the end of period t if either
 - (a) The responder rejected and $z_t > \alpha$.
 - (b) $s_t > s$ and the responder accepted.

Thus, in Phase 1 the claimant demands x and the responder accepts iff the signal s_t is sufficiently favorable; in Phase 2 the players play a trivial equilibrium; and play switches from Phase 1 to Phase 2 if either the responder rejects and the public randomization realization z_t is sufficiently high, or the responder deviates by accepting on a signal where he was supposed to reject. Note that both players' strategies are public.

Trigger strategies model the “trip wires” or “plate glass windows” that have long been thought to be essential for deterrence (Schelling, 1966, p. 47). A key feature of these strategies is that *both* players are disciplined by the threat of switching to a trivial equilibrium: the claimant because aggressive claims increase the probability that the responder rejects and the triggering event described in 3(a) occurs; and the responder because accepting after an unfavorable signal (which is tempting when $x < 1$) causes the triggering event described in 3(b) to occur.²²

Theorem 2 *Assume signals are public and unboundedly informative. For all $x \in (0, 1]$ and all $\eta > 0$, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there exists a trigger equilibrium (and thus a PPE) where the claimant demands x in Phase 1 and expected payoffs are within η of $(x, 1 - x)$.*

A couple technical remarks: First, the role of public randomization is to reduce the value destruction that results when punishment is triggered “accidentally.” Public randomization

²²Schelling (1966, p. 56) elegantly (and humorously) described the logic of the latter triggering event in the context of the Cold War: “Our deterrence rests on Soviet expectations. This, I suppose, is the ultimate reason why we have to defend California—aside from whether or not Easterners want to. There is no way to let California go to the Soviets and make them believe nevertheless that Oregon and Washington, Florida and Maine, and eventually Chevy Chase and Cambridge cannot be had under the same principle.”

could be dispensed with by making the transition to the trivial equilibrium temporary rather than permanent: from the players' perspectives, permanently transitioning to the trivial equilibrium with probability $1 - \alpha$ is equivalent to deterministically transitioning to the trivial equilibrium for T periods, where T satisfies $\delta^T = \alpha$. Second, the folk theorem for repeated games with imperfect public monitoring (Fudenberg, Levine, and Maskin, 1994) suggests that the unbounded informativeness condition can be relaxed, at the cost of using more complicated strategies. However, I am not aware of a general folk theorem that formally covers the current model, which is a repeated extensive-form game with continuous actions.

It is also worth noting that the feature that aggressive claims are deterred by the threat of future punishment in addition to current-period rejection is *necessary* for supporting an equilibrium payoff close to 1 for the responder. To see this, suppose that the responder uses a strategy that is *stationary on-path*, meaning that the responder uses the same threshold s^* in every period along the equilibrium path. (In other words, the strategy is of the trigger form but without contingency 3(a), or equivalently with $\alpha = 1$.²³) By the full support assumption, the claimant's continuation payoff remains constant even if she makes an off-path demand, so her problem is $\max_{\tilde{x} \in X} \tilde{x}F(s^*|\tilde{x})$, with first-order condition

$$\frac{F_x(s^*|x)}{F(s^*|x)} = -\frac{1}{x}.$$

To support an equilibrium payoff close to 1 for the responder, this condition must hold for some x close to 0. As $x \rightarrow 0$, this condition can hold only if $F_x(s^*|x)/F(s^*|x) \rightarrow -\infty$, and hence $s^* \rightarrow \underline{s}$ (since $F_x(s|x)/F(s|x)$ is increasing in s , as shown in Lemma 4 in the Appendix). But this in turn implies that $F(s^*|x) \rightarrow 0$, which contradicts with the responder's payoff being close to 1. Intuitively, if $x \approx 0$ and the responder's strategy is stationary on-path, the claimant loses very little when her equilibrium demand is rejected, so she can be deterred from marginally increasing the demand only if the equilibrium demand is rejected with high probability. In contrast, when the responder is patient, the threat of future punishment is severe even when x is small, so more aggressive demands can be deterred even when the equilibrium demand is accepted with high probability.

²³In particular, play still switches to a punishment phase if the responder deviates by accepting after an unfavorable signal.

3.2 Private Signals

In contrast to the situation with public signals, deterrence fails in any PPE when the signal is private.

Theorem 3 *Assume signals are observed only by the responder. In any PPE, conditional on reaching any public history h^t , either the claimant demands $x = 1$ with probability 1 or the responder rejects after every signal with probability 1 (or both).*

Thus, the possible on-path outcomes of a PPE in the repeated bargaining game with private signals are exactly the same as those in the one-shot game: either the claimant demands the entire pie and the responder accepts, or the responder rejects.

To see the intuition, suppose the responder accepts with positive probability in period t , conditional on reaching public history h^t . By a similar argument as in the one-shot game, this implies that the claimant's period- t demand is deterministic (conditional on reaching h^t). If the responder strictly prefers to accept this demand (taking into account any impact of his decision on his continuation payoff), and hence accepts with probability 1, he continues to accept with probability 1 if the claimant increases her demand. But then it is strictly optimal for the claimant to increase her demand, because, in a public equilibrium, the claimant's continuation payoff is independent of her period- t demand, conditional on the responder's period- t accept/reject decision. Therefore, the responder must be indifferent between accept and reject at any history where he accepts with positive probability, and hence his payoff is the same as if he always rejects. It follows that the responder accepts with positive probability only when the claimant demands 1.²⁴

While Theorem 3 shows that, with private signals, repeated game effects cannot deter aggressive demands (at least with public strategies), they can reduce disagreement. In particular, note that if $x = 1$ then the trigger strategy profile defined in Section 3.1 remains a well-defined public strategy profile when the signal is private, because the claimant's demand is the same in both phases, so the claimant does not need to observe the signal to play her strategy. The following result is therefore an immediate corollary of Theorem 2.

²⁴Theorem 3 is in the spirit of Matsushima's (1991) anti-folk theorem for repeated games with conditionally independent private monitoring, but my model does not satisfy his assumptions (for example, the responder's accept/reject decision is a public signal in my model) and the argument is different.

Corollary 1 *Assume signals are observed only by the responder and are unboundedly informative. For all $\eta > 0$, there exists $\bar{\delta} < 1$ such that, for every $\delta > \bar{\delta}$, there exists a trigger equilibrium where the claimant demands $x = 1$ in Phase 1 and the claimant’s expected payoff is greater than $1 - \eta$.*

Theorems 2 and 3 formalize a sense in which deterrence is qualitatively more difficult with private signals. These results have the practical implication that an actor who wants to deter aggressive claims can benefit by committing to publicize its information. For example, a well-known obstacle to deterrence in cyber-space is that firms and governments often prefer to conceal the fact that they have been hacked, rather than revealing the hack and facing the difficult question of how to respond (Sanger, 2019). Greater transparency in these situations can help facilitate deterrence.

I conclude this section by discussing the robustness of the impossibility of deterrence under private signals to allowing richer information structures and more complex, non-public equilibria. Realistically, the claimant should observe a private signal r of the responder’s signal s ; however, if we maintain the restriction to PPE, the claimant will not use this signal in equilibrium, and Theorem 3 remains valid. Allowing non-public equilibria complicates matters significantly, however (whether or not the claimant observes a signal of s). Probably the simplest equilibria that could secure a positive payoff for the responder are ones where the claimant continues to use a public strategy, but the responder uses a non-public strategy that conditions on past signal realizations. For example, if M is not too large one could imagine an equilibrium where in odd periods the claimant demands $x = 0.5$ and the responder always accepts, while in even periods the claimant demands $x = 1$ and the responder accepts iff $s_t \leq s^*(s_{t-1})$, where $s^*(\cdot)$ is a decreasing function. While these equilibria may be relatively easy to construct, they are clearly not robust to making the responder’s signal slightly payoff-relevant in the manner discussed in Section 2.3, and thus may not be very realistic.²⁵ Finally, it may be possible to construct more robust equilibria where both parties use non-public strategies, and in particular the claimant uses a mixed strategy where her current claim is conditioned on the realizations of her past private randomizations. It is not clear if such

²⁵This critique is closely related to the non-purifiability of belief-free equilibria in private-monitoring repeated games (e.g., Bhaskar, 1998).

equilibria can actually be constructed; if they can, they would likely be quite complicated, and possibly unrealistic.²⁶ I leave this as an open question.

3.3 Monotone Claims, Irreversible Rejection, and Salami Tactics

In some applications, it seems more realistic to assume that x_t must be monotone and/or that rejecting a claim is irreversible (in contrast to the repeated game model considered so far, where the claim and the accept/reject decision can be freely changed from period to period). The resulting model also captures the intuitive notion of “salami tactics,” where the claimant gradually claims resources and the responder decides whether and when to initiate a conflict.

As compared to the repeated game model, constraining x_t to be monotone and/or making rejection irreversible shrinks the players’ strategy spaces. Therefore, any equilibrium in the repeated game model with the property that x_t is monotone and rejection is “permanent” (meaning that whenever the responder rejects in period t , he also rejects in all future periods with probability 1) remains an equilibrium in the constrained model. Note that the trigger strategy profiles used in Section 3.1 involve monotone x_t but “impermanent” rejection, so these profiles remain equilibria if x_t is constrained to be monotone, but not if rejection is also constrained to be permanent. In this case, the following modification can be used.

Definition 2 *A modified trigger strategy profile is characterized by numbers $x \in (0, 1]$, $s \in S$, and $\alpha \in [0, 1]$, and is defined as follows:*

1. *In Phase 1, the claimant demands x and the responder accepts after every signal.*
2. *In Phase 2, the claimant demands 1 and the responder rejects after every signal.*
3. *Play begins in Phase 1, and permanently transitions to Phase 2 at the end of period t if $s_t > s$ and $z_t > \alpha$.*

²⁶Such equilibria would be “belief-based” in the sense of Sekiguchi (1997) and Bhaskar and Obara (2002). This type of equilibrium is notoriously difficult to analyze.

Thus, the responder now accepts after every signal, and the claimant is deterred from increasing her demand entirely through the threat of switching to permanent rejection.²⁷

Proposition 5 *If one replaces “trigger equilibrium” with “modified trigger equilibrium,” Theorem 2 remains valid under the monotone claims and/or irreversible rejection restrictions. Theorem 3 also remains valid under the monotone claims and/or irreversible rejection restrictions.*

I omit the proof of Proposition 5. For the first part, the proof is a variation on the proof of Theorem 2, with slightly different formulas. For the second part, the proof of Theorem 3 can be read verbatim.

The model’s predictions regarding the players’ payoffs are thus unaffected by requiring that claims are monotone and/or rejection is irreversible. However, if rejection is irreversible, one obtains a novel prediction regarding the long-run outcome of the game: ultimately, deterrence breaks down and conflict prevails. This prediction applies for any Nash equilibrium, regardless of whether signals are public or private, and regardless of whether claims are restricted to be monotone.

Proposition 6 *Fix any NE of the repeated bargaining game with irreversible rejection. Letting $q_t = \Pr(a_\tau = \text{accept} \forall \tau \leq t)$ denote the probability that the responder accepts in each of the first t periods, we have $q_t \rightarrow 0$.*

The intuition is that if $q_t \rightarrow 0$ then there exists a period T beyond which the equilibrium rejection probability is very small. By the full support assumption, this implies that the rejection probability starting from period T is also very small if the claimant deviates by demanding $x_t = M$ in every period $t \geq T$. But then this deviation is profitable.

4 Alternating Unobserved Offers

The baseline model of Section 2 captures asymmetric situations where one party claims resources and the other party can either acquiesce or start a conflict. This setup fits some

²⁷Modified trigger strategies could also have been used in Section 3.1. I prefer the original definition of trigger strategies because they seem more descriptive, and also because the claimant’s strategy continues to satisfy the definition of a public strategy in the private-signals model.

of my motivating applications well, such as the politician–social media platform example, or other regulation or agency problems. Other applications, such as international relations, are more symmetric, and in these applications it may be better to view the parties as taking turns claiming resources. A natural model of these more symmetric situations is Rubinstein’s (1982) alternating-offers bargaining model with imperfectly-observed offers. This is simply the alternating-offers version of the baseline model.

Formally, the game proceeds in discrete periods, $t = 1, 2, \dots$. In odd periods, the game is the one-shot bargaining model of Section 2, where player 1 makes a demand $x_1 \in [-M, M]$ and then player 2 observes a signal $s_1 \in [\underline{s}_1, \bar{s}_1]$, distributed $F_1(s_1|x_1)$, before accepting or rejecting. In even periods, the roles are reversed: player 2 makes a demand $x_2 \in [-M, M]$ and then player 1 observes a signal $s_2 \in [\underline{s}_2, \bar{s}_2]$, distributed $F_2(s_2|x_2)$, and then accepts or rejects. I assume for simplicity that the players have the same discount factor $\delta \in [0, 1)$, so if player $j \neq i$ accepts player i ’s demand x_i in period t , the game ends with payoffs of $\delta^{t-1}x_i$ for player i and $\delta^{t-1}(1 - x_i)$ for player j .²⁸ When an offer is rejected, play continues on to the next period. The signal distributions F_1 and F_2 can differ from each other, but both are assumed to satisfy Assumption 1. Observe that if $\delta = 0$, we recover the one-shot model of Section 2; and if F_1 and F_2 are degenerate, we recover Rubinstein’s model.

With perfectly-observed offers, Rubinstein showed that this model has a unique subgame-perfect equilibrium. With imperfectly-observed offers, there are already multiple equilibria in ultimatum bargaining (as shown in Theorem 1), so such a strong uniqueness result cannot hold. For example, there is a PPE in the alternating-offers game where the Pareto-optimal equilibrium of the one-shot bargaining game is played in period 1, and this period-1 play is supported by having the players permanently switch to a trivial (no-trade) equilibrium starting in period 2 if player 2 rejects in period 1. As shown in Proposition 2, player 1’s payoff in this equilibrium converges to 1 as observation noise vanishes. Similarly, there is also a PPE where player 2 always rejects in period 1 and the Pareto-optimal one-shot equilibrium is played in period 2, so now it is player 2’s payoff that converges to 1 as noise vanishes.

Despite this multiplicity, there is a natural class of equilibria to focus on, which generalizes the Pareto-optimal equilibrium of the one-shot game while imposing a stationarity

²⁸Throughout this section, i and j indicate distinct players.

assumption. These are *stationary cutoff equilibria (SCE)*, where for each $i \in \{1, 2\}$ there exist numbers $x_i \in [-M, M]$ and $s_i^* > \underline{s}_i$ such that player i claims x_i whenever it is her turn to make a claim, and player j accepts iff $s_i \leq s_i^*$ whenever it is his turn to respond. Note that the condition that $s_i^* > \underline{s}_i$ for each i together with the full-support assumption implies that each player's claim is accepted with positive probability in every period.

I will show that, under the following additional assumption on F_1 and F_2 , (i) an SCE exists, (ii) the set of SCE is ordered, where equilibria that are better for player 1 are worse for player 2, and (iii) the extremal SCE (i.e., the best and worst SCE for each player) exhibit interesting comparative statics.

Assumption 2 For each $i \in \{1, 2\}$, the signal distribution F_i satisfies

1. *Additive noise:* There exists a distribution G_i and a parameter θ_i such that $F_i(s_i|x_i) = G_i((s_i - x_i)/\theta_i)$ for all s_i, x_i .
2. *Informativeness:* Denoting the support of G_i by $[\underline{z}_i, \bar{z}_i]$ (where $\underline{z}_i \in \mathbb{R} \cup \{-\infty\}$ and $\bar{z}_i \in \mathbb{R} \cup \{+\infty\}$) and denoting the density of G_i by g_i , we have $\lim_{z \rightarrow \underline{z}_i} g_i(z)/G_i(z) > \theta_i/(1 - \delta)$ and $\lim_{z \rightarrow \bar{z}_i} g_i(z) < \theta_i/(1 - \delta^2)$.

The additive noise assumption simplifies the analysis by ensuring that the sensitivity of the acceptance probability to i 's demand, $F_{i,x}(s_i^*|x_i) := \partial F_i(s_i^*|x_i)/\partial x_i$, is pinned down by the acceptance probability $F_i(s_i^*|x_i)$ itself, rather than depending separately on s_i^* and x_i . This property will be useful in relating the first-order condition for x_i to be an optimal demand for player i and the Rubinstein-type equation that determines x_i as a function of the acceptance probabilities (equation (3) below). The additive noise assumption also gives a parameter θ_i for measuring the noisiness of i 's claim. The informativeness assumption will be used to establish existence of an SCE, similarly to how the earlier informativeness condition was used in Section 2.

In what follows, I parameterize an SCE by the equilibrium demands (x_1, x_2) and the equilibrium rejection probabilities (p_1, p_2) , where $p_i = 1 - F_i(s_i^*|x_i)$ for $i \in \{1, 2\}$. I also denote player i 's equilibrium expected utility by U_i . Since a player is always indifferent between accepting and rejecting her opponent's offer in an SCE, and since player 1 makes

the first offer, we have

$$U_1 = (1 - p_1)x_1 + \delta p_1(1 - x_2) \quad \text{and} \quad U_2 = 1 - x_1. \quad (2)$$

My main result for the alternating unobserved offers model is as follows.

Theorem 4 *With alternating unobserved offers, the following hold:*

1. *An SCE exists.*
2. *The set of SCE is completely ordered: for any two SCE (x_1, x_2, p_1, p_2) and (x'_1, x'_2, p'_1, p'_2) with payoffs (U_1, U_2) and (U'_1, U'_2) , if $x_i \geq x'_i$ for some $i \in \{1, 2\}$, then $x_j \leq x'_j$, $p_i \leq p'_i$, $p_j \geq p'_j$, $U_i \geq U'_i$, and $U_j \leq U'_j$.*
3. *For each $i \in \{1, 2\}$, the smallest and largest SCE values of x_i and U_i are both decreasing in θ_i and increasing in θ_j , and the smallest and largest SCE values of p_i are both increasing in θ_i and decreasing in θ_j .*

Part 2 of the theorem says that the players' preferences over distinct SCE are opposed: better SCE for player i are worse for player j . Moreover, in a better SCE, a player demands more and has her demand accepted more often; while her opponent demands less and has his demand accepted less often. Part 3 of the theorem says that, at least for extremal SCE, a player benefits when her own demand is observed more precisely, as well as when her opponent's demand is observed less precisely. That is, the ability to make clear demands (and the opponent's inability to do so) is an advantage in bargaining.²⁹

The logic of Theorem 4 is that an SCE is determined by four equations, which jointly determine the four equilibrium variables (x_1, x_2, p_1, p_2) . The equations are, for each player $i \in \{1, 2\}$, the condition that player i is indifferent between accepting and rejecting her opponent's demand, and the condition that player i 's own demand is optimal given the cutoff signal s_i^* used by the opponent. The former equations are

$$1 - x_j = \delta(1 - p_i)x_i + \delta^2 p_i(1 - x_j) \quad \text{for } i \in \{1, 2\}, j \neq i.$$

²⁹An open question is whether some natural assumptions on (F_1, F_2) might rule out multiple SCE. For the range of parameters in Figures 2–4 below, it can be verified numerically that the SCE is unique.

The explanation for these equations is familiar from Rubinstein-type bargaining models. For each $i \in \{1, 2\}$, they say that player i is indifferent between accepting player j 's demand (which gives payoff $1 - x_j$) and rejecting it and returning with her own demand in the next period (which gives payoff $\delta(1 - p_i)x_i + \delta^2 p_i(1 - x_j)$), noting that player i 's next-period demand is itself accepted with probability $1 - p_i$. The two equations can be solved for (x_i, x_j) as a function of (p_i, p_j) , to obtain

$$x_i = \frac{(1 + \delta p_j)(1 - \delta^2 p_i)}{(1 + \delta)(1 - \delta^2 p_i p_j)} \quad \text{for } i \in \{1, 2\}, j \neq i. \quad (3)$$

Equation (3) characterizes the unique subgame-perfect equilibrium of a Rubinstein-type bargaining model where each of player i 's offers is “lost” with exogenous probability p_i . Note that x_i is decreasing in p_i and increasing in p_j . That is, in Rubinstein-type bargaining, a player benefits when her opponent's offers go missing with higher probability, and she suffers when her own offers go missing. This easy observation is a “baby version” of part 3 of Theorem 4. The actual result also accounts for the fact that p_i and p_j are endogenous.³⁰

The remaining equilibrium conditions require that each player i 's demand is optimal given the opponent's cutoff signal s_i^* : that is,

$$x_i = \operatorname{argmax}_{\tilde{x}_i} F_i(s_i^* | \tilde{x}_i) \tilde{x}_i + (1 - F_i(s_i^* | \tilde{x}_i)) \delta(1 - x_j) \quad \text{for } i \in \{1, 2\}, j \neq i.$$

This is similar to the claimant's problem in the baseline model, except that now player i 's continuation payoff when her demand is rejected equals $\delta(1 - x_j)$, rather than 0. The first-order condition (which, as in the baseline model, is necessary and sufficient) is

$$F_i(s_i^* | x_i) + (x_i - \delta(1 - x_j)) F_{i,x}(s_i^* | x_i) = 0.$$

Substituting for x_i and x_j using the Rubinstein equations, (3), and simplifying, we have

$$x_i - \delta(1 - x_j) = \frac{(1 - \delta)(1 + \delta p_j)}{1 - \delta^2 p_i p_j} > 0,$$

³⁰Rubinstein-type bargaining with an exogenous chance of losing offers fits into the general framework of Binmore (1987) or Merlo and Wilson (1993), so the baby result that x_i is decreasing (increasing) in *exogenous* p_i (p_j) is not entirely novel.

so the first-order condition can be written as

$$\frac{(1 - \delta)(1 + \delta p_j)}{1 - \delta^2 p_i p_j} \left(\frac{-F_{i,x}(s_i^*|x_i)}{F_i(s_i^*|x_i)} \right) = 1.$$

Since $F_i(s_i^*|x_i) = 1 - p_i$ and $-F_{i,x}(s_i^*|x_i) = g_i(G_i^{-1}(1 - p_i))/\theta_i$ by the additive noise assumption, the first-order condition further simplifies to

$$\frac{(1 - \delta)(1 + \delta p_j) g_i(G_i^{-1}(1 - p_i))}{(1 - \delta^2 p_i p_j)(1 - p_i)} = \theta_i. \quad (4)$$

The proof is completed by showing that equation (4) defines a continuous, downward-sloping curve C_i in (p_i, p_j) -space (which follows because $g_i(z)/G_i(z)$ is decreasing in z , by MLRP), and that the curves C_i and C_j intersect (which follows from the informativeness assumption). This observation implies that an SCE exists, and together with the Rubinstein equations it implies that the set of SCE has the desired order structure. Finally, increasing θ_i shifts the curve C_i upwards and therefore increases the smallest and greatest SCE values of p_i , which together with the order structure implies part 3 of theorem.

Figures 2–4 illustrate Theorem 4 in the normal-noise case, $G_i = \mathcal{N}(0, \theta_i^2)$ for $i \in \{1, 2\}$. The figures graph the SCE demands (x_1, x_2) , rejection probabilities (p_1, p_2) , and expected payoffs (U_1, U_2) as a function of $\theta_1^2 \in [0.5, 1.5]$, when θ_2^2 is fixed at 1 and δ is fixed at 0.9 (with the player 1 variables in red). The comparative statics of Theorem 4 are clearly visible. Another interesting observation is that the acceptance probabilities are lower with alternating offers than with an ultimatum offer. For example, the acceptance probability in ultimatum bargaining with $\theta^2 = 1$ is 38%, but the acceptance probability in alternating-offers bargaining with $\theta_1^2 = \theta_2^2 = 1$ and $\delta = 0.9$ is less than 10%. Thus, in this parametric example, the ability to make counter-offers not only moderates the opponent’s demand (as in Rubinstein’s model), but also increases the probability of rejecting it (unlike in Rubinstein). This seems like a realistic prediction of the model.




Figure 2: Demands




Figure 3: Rejection Probabilities




Figure 4: Expected Payoffs

5 Conclusion

This paper has developed a simple theory of disagreement in bargaining, based on the idea that a party's demand or claim—which in reality often corresponds to a substantive action, rather than a formal term in a proposed contract—is imperfectly observed. Imperfectly observed demands cannot always be accepted, because if they were, the claimant would ask for more. The resulting bargaining model is tractable and has many intuitive features: for example, demands are more likely to be accepted when they are observed more precisely. In repeated bargaining, trigger strategies can deter aggressive claims if the signal of the claim is publicly observed, but not if it is observed only by the responding party. In alternating-offers bargaining, the ability to make clear demands is an advantage, as is being known for an inability to accurately assess others' demands.

Rational-choice theorists traditionally view bargaining failure as a puzzle in need of explanation. This perspective is rooted in the assumption that bargainers can make perfectly clear offers, and thus can avoid disagreement by moderating their demands. This assumption is often a useful simplification but it is probably not very realistic, especially when an “offer” is a concrete state of affairs rather than a number of dollars or some other unambiguously measurable quantity. The current paper is a preliminary study of the consequences of relaxing this assumption. Among many further questions that could be asked about unobserved-offers bargaining, a particularly interesting one is whether empirical evidence supports the proposed connection between noisy monitoring of claims and conflict risk.

A Appendix: Omitted Proofs

A.1 Proof of Theorem 1

It is obvious that any strategy profile where $\mathbb{E}[x|s] \geq 1$ for all $s \in S$ and the responder always rejects is a trivial NE; and conversely that in every trivial NE it must be the case that $\mathbb{E}[x|s] \geq 1$ for all $s \in S$. I show that every non-trivial NE takes the prescribed form.

Lemma 1 *In any NE, the responder accepts with probability strictly less than 1.*

Proof. If the responder accepts with equilibrium probability 1, the full support assumption implies that he accepts with probability 1 after every signal. The claimant's unique best response to this strategy is demanding $x = M$. But the responder's best response to $x = M$ is to always reject. ■

Lemma 2 *The claimant uses a pure strategy in every non-trivial NE.*

Proof. Suppose toward a contradiction that the claimant uses a mixed strategy. By a standard property of MLRP (Milgrom, 1981, Proposition 1), for any non-degenerate prior distribution G of x , if the conditional distribution $F(s|x)$ satisfies full support and strict MLRP, then $G(x|s) > G(x|s')$ for all $s < s'$ and all x such that $0 < G(x) < 1$. Therefore, $\mathbb{E}[x|s] < \mathbb{E}[x|s']$ for all $s < s'$. The latter property implies that the responder must use a cutoff strategy: accept iff $s \leq s^*$, for some $s^* \in S$.³¹ Moreover, by full support and the hypothesis that the responder accepts with positive probability (as the equilibrium is non-trivial), $s^* > \underline{s}$. Also, since the responder cannot accept with probability 1 (by Lemma 1), $s^* < \bar{s}$.

A contradiction can now be derived by showing that the claimant's best response to any cutoff strategy satisfying $\underline{s} < s^* < \bar{s}$ is pure. Note that the claimant's expected utility from demanding x equals $xF(s^*|x)$. Strict MLRP, $\underline{s} < s^* < \bar{s}$, and full support imply that $F(s^*|x) > F(s^*|x')$ for all $x < x'$, and hence $F_x(s^*|x) < 0$ for all x . Finally, if $F_x(s^*|x) < 0$ for all x and F is log-concave in x , then $xF(s^*|x)$ has a unique maximizer. Indeed, if x is a

³¹Recall that I identified strategies that differ on a measure-0 set, so the responder's behavior at $s = s^*$ is immaterial.

critical point of $xF(s^*|x)$, then $xF_x(s^*|x) + F(s^*|x) = 0$, and hence the second derivative of $xF(s^*|x)$ at the critical point equals

$$xF_{xx}(s^*|x) + 2F_x(s^*|x) = -\frac{F(s^*|x)F_{xx}(s^*|x)}{F_x(s^*|x)} + 2F_x(s^*|x).$$

Since $F_x(s^*|x) < 0$ and $F(s^*|x)F_{xx}(s^*|x) \leq (F_x(s^*|x))^2$ (by log-concavity in x), the second derivative is strictly negative. Hence, every local extremum of $xF(s^*|x)$ is a strict local maximum, so $xF(s^*|x)$ has a unique global maximizer. ■

Lemmas 1 and 2 imply that the claimant demands $x = 1$ in any non-trivial NE: if the claimant demanded $x < 1$, the responder would always accept, contrary to Lemma 1; and if the claimant demanded $x > 1$, the responder would always reject, so the equilibrium could not be non-trivial. Therefore, $x = 1$ must be an optimal demand for the claimant: that is, $1 \in \operatorname{argmax}_{x \in X} x \int_{s \in S} \sigma(s) f(s|x) ds$. This completes the proof of the first part of Theorem 1.

I now show that, if a non-trivial NE exists, the probability that the responder accepts is maximized at an equilibrium where the responder accepts iff $s \leq s^*$, where $s^* > \underline{s}$ satisfies (1). To see this, note that the probability that the responder accepts equals $\int \sigma(s) f(s|x=1) ds$, while a necessary condition for $x = 1$ to be an optimal demand for the claimant (i.e., the first-order condition of the claimant's problem $\max_x x \int \sigma(s) f(s|x) ds$ at $x = 1$) is $\int \sigma(s) (f(s|x=1) + f_x(s|x=1)) ds = 0$. By a standard Lagrangian argument, maximizing the probability that the responder accepts over strategies $\sigma(s)$ subject to the first-order condition implies that $\sigma(s) = 1$ if $f_x(s|x=1)/f(s|x=1)$ is below some threshold $\lambda > 0$, and that $\sigma(s) = 0$ if $f_x(s|x=1)/f(s|x=1) > \lambda$. Since $f_x(s|x=1)/f(s|x=1)$ is strictly increasing in s under full support and MLRP (Milgrom, 1981; Proposition 5), this implies that the probability that the responder accepts is maximized by an equilibrium where the responder accepts iff $s \leq s^*$, for some s^* . Finally, when the responder's strategy takes this form, the claimant's first-order condition at $x = 1$ is (1).

It remains to show that a non-trivial equilibrium exists iff the informativeness condition holds. Since a non-trivial equilibrium exists iff such an equilibrium exists where the responder accepts iff $s \leq s^*$ for some $s^* > \underline{s}$, I restrict attention to equilibria of this cutoff form. As

was shown in the proof of Lemma 2, the claimant's problem has a unique local maximizer when the responder uses a cutoff strategy. Therefore, there exists a non-trivial equilibrium of the cutoff form iff there exists $s^* > \underline{s}$ that satisfies (1). The following lemma therefore completes the proof.

Lemma 3 *There exists $s^* > \underline{s}$ that satisfies (1) iff the informativeness condition holds*

I first record a simple implication of full support and MLRP.

Lemma 4 *For any x and $s' > s$, we have $F_x(s'|x)/F(s'|x) > F_x(s|x)/F(s|x)$.*

Proof. Suppressing the conditioning event to ease notation, we have

$$\begin{aligned} \frac{d F_x(s)}{ds F(s)} &= \frac{1}{F(s)^2} (F(s) f_x(s) - f(s) F_x(s)) \\ &= \frac{1}{F(s)^2} \left(\int_{-\infty}^s f(\tilde{s}) f_x(s) d\tilde{s} - \int_{-\infty}^s f(s) f_x(\tilde{s}) d\tilde{s} \right) \\ &= \frac{1}{F(s)^2} \int_{-\infty}^s (f(\tilde{s}) f_x(s) - f_x(\tilde{s}) f(s)) d\tilde{s}, \end{aligned}$$

which is strictly positive because $f(\tilde{s}) f_x(s) > f_x(\tilde{s}) f(s)$ for all $\tilde{s} < s$ (Milgrom, 1981, Proposition 5). ■

Proof of Lemma 3. Let

$$h(x, s) = F(s|x) + xF_x(s|x). \tag{5}$$

Note that (1) holds iff $h(1, s^*) = 0$. Since F and F_x are continuous in s , there exists $s^* > \underline{s}$ such that $h(1, s^*) = 0$ iff there exist $s, s' > \underline{s}$ such that $h(1, s) \leq 0 \leq h(1, s')$.

As seen in the proof of Lemma 2, for every s , $h(x, s)$ is single-crossing from above in x : if $h(x, s) = 0$ then $h(x', s) > 0$ for all $x' < x$ and $h(x', s) < 0$ for all $x' > x$. Since F is a cdf and $M > 1$, there exists $s' > \underline{s}$ such that $F(s'|x = M) > 1/M$. Since $MF(s'|x = M) > 1$, we have $\operatorname{argmax}_x xF(s'|x) > 1$. Since $h(x, s') = (d/dx) xF(s'|x)$ is single-crossing from above in x , we have $h(1, s') > 0$.

It remains to show that there exists $s > \underline{s}$ such that $h(1, s) \leq 0$ iff the informativeness condition holds. For all $s > \underline{s}$, we have $F(s|x=1) > 0$ by full support, so $h(1, s) \leq 0$ iff

$$\frac{F_x(s|x=1)}{F(s|x=1)} \leq -1. \quad (6)$$

By Lemma 4, $F_x(s|x=1)/F(s|x=1)$ is strictly increasing in s . Hence, there exists $s > \underline{s}$ such that $h(1, s) \leq 0$ iff

$$\lim_{s \rightarrow \underline{s}} \frac{F_x(s|x=1)}{F(s|x=1)} < -1.$$

If $F_x(s|x=1)$ does not converge to 0 as $s \rightarrow \underline{s}$, then the left hand side of this inequality equals $-\infty$, and the informativeness condition holds. If instead $\lim_{s \rightarrow \underline{s}} F_x(s|x=1) = 0$, then by l'Hopital's rule, we have

$$\lim_{s \rightarrow \underline{s}} \frac{F_x(s|x=1)}{F(s|x=1)} = \lim_{s \rightarrow \underline{s}} \frac{f_x(s|x=1)}{f(s|x=1)},$$

which is less than -1 iff the informativeness condition holds. ■

A.2 Proof of Proposition 1

I show that the claimant always demands $x = 1$ in any non-trivial equilibrium. Viewing $\sigma(s)$ as the probability that the responder accepts conditional on signal realization s , the rest of the proof follows that of Theorem 1.

Fix a non-trivial equilibrium. As in the proof of Lemma 2, if the distribution of x conditional on message m is non-degenerate, the responder accepts iff $s \leq s^*(m)$, for some message-contingent cutoff $s^*(m)$. Furthermore, if the distribution of x conditional on message m is degenerate on some $x \neq 1$, the responder accepts iff $s \leq s^*(m) \in \{\underline{s}, \bar{s}\}$ (i.e., he always rejects, or always accepts). Hence, for every message m , either the distribution of x conditional on m is degenerate on $x = 1$, or the responder accepts iff $s \leq s^*(m)$. Let $\bar{s}^* = \sup_m s^*(m)$. Since the signal has full support, it is suboptimal for the claimant to send a message m that induces a cutoff strategy with $s^*(m) < \bar{s}^*$. Hence, the claimant only sends messages m such that either $s^*(m) = \bar{s}^*$ or the distribution of x conditional on m is degenerate on $x = 1$. Call the former set of messages M_1 and the latter set M_2 . Since the

result is immediate if the probability that the claimant sends a message in M_1 is 0, assume that this probability is strictly positive.

As in the proof of Lemma 2, there is a unique maximizer of $x F(\bar{s}^*|x)$: call it x^* . Therefore, either the claimant demands x^* and sends a message in M_1 , or she demands 1 and sends a message in M_2 . By the same reasoning as in Lemma 1, the responder must accept with probability strictly less than 1 after receiving a message in M_1 : therefore, $x^* \geq 1$. Since the responder accepts with positive probability after receiving a message in M_2 (since the equilibrium is non-trivial), he must also accept with positive probability after receiving a message in M_1 , or else the responder will never send such a message (contrary to our hypothesis): therefore, $x^* \leq 1$. Hence, $x^* = 1$, so the claimant always demands $x = 1$.

A.3 Proof of Proposition 2

Consider a convergent subsequence and suppose that $s^{n,*} \rightarrow \hat{s} \neq 1$. Suppose that $\hat{s} > 1$. By $1 \in \operatorname{argmax}_{x \in X} x F^n(s^{n,*}|x)$, we have

$$F^n(s^{n,*}|x=1) \geq \frac{1 + \min\{s^{n,*}, M\}}{2} F^n\left(s^{n,*}|x = \frac{1 + \min\{s^{n,*}, M\}}{2}\right).$$

Since $s^{n,*} \rightarrow \hat{s} > 1$ and $\sup_{s,x} |F^n(s|x) - \mathbf{1}\{s \geq x\}| \rightarrow 0$, we have

$$\frac{1 + \min\{s^{n,*}, M\}}{2} F^n\left(s^{n,*}|x = \frac{1 + \min\{s^{n,*}, M\}}{2}\right) \rightarrow \frac{1 + \min\{\hat{s}, M\}}{2} > 1,$$

which is a contradiction since $F^n(s^{n,*}|x=1) \leq 1$.

Similarly, suppose that $\hat{s} < 1$. Since $f_x^n(0|1)/f^n(0|1) < -1$ for all n , the fact that $F_x(s|x=1)/F(s|x=1)$ is strictly increasing in s (by Lemma 4), together with (1), implies that $\hat{s} > 0$. By $1 \in \operatorname{argmax}_{x \in X} x F^n(s^{n,*}|x)$, we have

$$F^n(s^{n,*}|x=1) \geq \frac{s^{n,*}}{2} F^n\left(s^{n,*}|x = \frac{s^{n,*}}{2}\right).$$

Since $s^{n,*} \rightarrow \hat{s} < 1$ and $\sup_{s,x} |F^n(s|x) - \mathbf{1}\{s \geq x\}| \rightarrow 0$, we have

$$F^n(s^{n,*}|x=1) \rightarrow 0 \text{ and } \frac{s^{n,*}}{2} F^n\left(s^{n,*}|x = \frac{s^{n,*}}{2}\right) \rightarrow \frac{\hat{s}}{2} > 0,$$

which is a contradiction. This shows that all convergent subsequences converge to $\hat{s} = 1$. As the same argument also contradicts the existence of a subsequence diverging to $-\infty$ or $+\infty$, the space of possible cutoff signals can be taken to be compact, so the original sequence also converges to $\hat{s} = 1$.

Finally, suppose that $\hat{s} = 1$ and $F^n(s^{n,*}|1) \rightarrow \alpha < 1$. By $1 \in \operatorname{argmax}_{x \in X} xF^n(s^{n,*}|x)$, we have

$$F^n(s^{n,*}|x=1) \geq \frac{1+\alpha}{2} F^n\left(s^{n,*}|x = \frac{1+\alpha}{2}\right).$$

Since $\alpha < 1$ and $\sup_{s,x} |F^n(s|x) - \mathbf{1}\{s \geq x\}| \rightarrow 0$, we have

$$F^n(s^{n,*}|x=1) \rightarrow \alpha \text{ and } \frac{1+\alpha}{2} F^n\left(s^{n,*}|x = \frac{1+\alpha}{2}\right) \rightarrow \frac{1+\alpha}{2} > \alpha,$$

which is a contradiction. Therefore, we also have $F^n(s^{n,*}|1) \rightarrow 1$.

A.4 Proof of Proposition 3

Since $S = \mathbb{R}$ and $\pi < 1$, for any strategy for the claimant, the responder accepts if s is a sufficiently large negative number. Therefore, every equilibrium is non-trivial. By the same argument as in the proof of Lemma 2, this implies that the claimant uses a pure strategy. Letting $x^*(\pi)$ denote the claimant's (pure) equilibrium demand, the responder's unique best response is to accept iff $s \leq s^*(\pi)$, where $s^*(\pi) = (1 - \pi x^*(\pi)) / (1 - \pi)$. Hence, the claimant's problem is $\max_{x \in X} xF(s^*(\pi)|x)$.

I claim that $x^*(\pi) \rightarrow 1$ as $\pi \rightarrow 1$. To see this, take a convergent subsequence and suppose that $x^*(\pi) \rightarrow \hat{x}$. Suppose further that $\hat{x} < 1$. This implies that $s^*(\pi) \rightarrow \infty$. But since F is a cdf and $M > 1$, for sufficiently large s the solution to $\max_{x \in X} xF(s|x)$ is greater than 1, which contradicts the hypothesis that $\operatorname{argmax}_{x \in X} xF(s^*(\pi)|x) < 1$ for π sufficiently close to 1.

Next, suppose that $\hat{x} > 1$, and therefore $s^*(\pi) \rightarrow -\infty$. Note that the first-order condition

of the claimant's problem is

$$\frac{F_x(s^*(\pi)|x^*(\pi))}{F(s^*(\pi)|x^*(\pi))} = -\frac{1}{x^*(\pi)}.$$

Since $s^*(\pi) \rightarrow -\infty$ and $F_x(s, x)/F(s, x)$ is continuous, the left-hand side of this equation converges to $\lim_{s \rightarrow -\infty} f_x(s|\hat{x})/f(s|\hat{x})$. By log-concavity (and l'Hôpital's rule), $\hat{x} > 1$, and the informativeness condition, we have

$$\lim_{s \rightarrow -\infty} \frac{f_x(s|\hat{x})}{f(s|\hat{x})} = \lim_{s \rightarrow -\infty} \frac{F_x(s|\hat{x})}{F(s|\hat{x})} \leq \lim_{s \rightarrow -\infty} \frac{F_x(s|1)}{F(s|1)} = \lim_{s \rightarrow -\infty} \frac{f_x(s|1)}{f(s|1)} < -1 < -\frac{1}{\hat{x}},$$

which is a contradiction. This shows that $\hat{x} = 1$. Since the convergent subsequence was arbitrarily chosen, compactness of X implies that the original sequence converges to 1.

I have established that $x^*(\pi) \rightarrow 1$ as $\pi \rightarrow 1$. Since $F(s, x)$ is continuous and $x^*(\pi) = \operatorname{argmax}_{x \in X} xF(s^*(\pi)|x)$, by the maximum theorem $s^*(\pi) \rightarrow \hat{s}$ such that $1 \in \operatorname{argmax}_{x \in X} xF(\hat{s}|x)$. This implies that $\hat{s} \rightarrow s^*$, given by (1).

A.5 Proof of Proposition 4

Note that $F(s^*|x=1) = G((s^* - 1)/\theta) = G(z^*)$, where $z^* := (s^* - 1)/\theta$. The claimant's first-order condition is

$$0 = F(s^*|1) + F_x(s^*|1) = G(z^*) - \frac{g(z^*)}{\theta}, \text{ or } \theta = \frac{g(z^*)}{G(z^*)}.$$

Totally differentiating the first-order condition with respect to θ , we have

$$1 = \frac{d}{d\theta} \theta = \frac{d}{d\theta} \frac{g(z^*)}{G(z^*)} = \frac{1}{G(z^*)^2} (G(z^*) g'(z^*) - g(z^*)^2) \frac{dz^*}{d\theta}.$$

At the same time, since F is log-concave, we have

$$F(s^*|x=1) F_{xx}(s^*|x=1) - F_x(s^*|x=1)^2 = \frac{1}{\theta^2} (G(z^*) g'(z^*) - g(z^*)^2) < 0.$$

Combining these equations, we have $dz^*/d\theta < 0$. Hence, $G(z^*) = F(s^*|x = 1)$ is likewise decreasing in θ .

A.6 Proof of Theorem 2

Fix $x \in (0, 1]$ and $\eta > 0$, and consider the trigger strategy profile with demand x , with thresholds s and α to be determined later. Let V denote the claimant's continuation payoff in Phase 1, and let $\beta = (1 - \alpha) / (1 - \delta)$. We have

$$\begin{aligned} V &= (1 - \delta) x F(s|x) + (1 - (1 - \alpha)(1 - F(s|x))) \delta V \\ &= (1 - \delta) x F(s|x) + (1 - \beta(1 - \delta)(1 - F(s|x))) \delta V \\ &= \frac{(1 - \delta) x F(s|x)}{1 - \delta(1 - \beta(1 - \delta)(1 - F(s|x)))} \\ &= \frac{x F(s|x)}{1 + \beta\delta(1 - F(s|x))}. \end{aligned}$$

The claimant's optimality condition is

$$x \in \operatorname{argmax}_{x'} (1 - \delta) x' F(s|x') + (1 - \beta(1 - \delta)(1 - F(s|x'))) \delta V,$$

with first-order condition

$$F(s|x) + x F_x(s|x) + \beta\delta F_x(s|x) V = 0.$$

We show that s and β can be chosen so that the first-order condition holds and $V > x - \eta$.

Let

$$g(x) = \sup_{s \in S} \frac{F_x(s|x)}{F(s|x)(1 - F(s|x))},$$

and, for any $\tilde{\eta} > 0$, let

$$\gamma(\tilde{\eta}) = \begin{cases} \tilde{\eta}/2 & \text{if } 1 + xg(x) \geq 0 \\ \min\left\{\tilde{\eta}, -\frac{1}{1+xg(x)}\right\}/2 & \text{if } 1 + xg(x) < 0 \end{cases}.$$

Let $s(\tilde{\eta})$ be the greatest solution to

$$\frac{F_x(s|x)}{F(s|x)(1-F(s|x))} = -\frac{1+\gamma(\tilde{\eta})}{\gamma(\tilde{\eta})x}. \quad (7)$$

This is well-defined because the left-hand side of (7) is continuous by assumption, and we have

$$\sup_{s \in S} \frac{F_x(s|x)}{F(s|x)(1-F(s|x))} = g(x) = -\frac{1 + \left(-\frac{1}{1+xg(x)}\right)}{\left(-\frac{1}{1+xg(x)}\right)x} > -\frac{1+\gamma(\tilde{\eta})}{\gamma(\tilde{\eta})x},$$

and

$$\inf_{s \in S} \frac{F_x(s|x)}{F(s|x)(1-F(s|x))} \leq \lim_{s \rightarrow \bar{s}} \frac{F_x(s|x)}{1-F(s|x)} = \lim_{s \rightarrow \bar{s}} -\frac{f_x(s|x)}{f(s|x)} < -\frac{1+\gamma(\tilde{\eta})}{\gamma(\tilde{\eta})x},$$

where the last inequality holds by unbounded informativeness. Moreover, note that $\lim_{\tilde{\eta} \rightarrow 0} \gamma(\tilde{\eta}) = 0$, and therefore $\lim_{\tilde{\eta} \rightarrow 0} s(\tilde{\eta}) = \bar{s}$, again by unbounded informativeness.

We now fix s . Rearranging (7) gives

$$\frac{\gamma(\tilde{\eta})}{1-F(s(\tilde{\eta})|x)} = -\frac{(1+\gamma(\tilde{\eta}))F(s(\tilde{\eta})|x)}{xF_x(s(\tilde{\eta})|x)}.$$

Since $\lim_{s \rightarrow \bar{s}} F_x(s|x) = 0$, because $F(s|x)$ is a cdf for every x , we have

$$\lim_{\tilde{\eta} \rightarrow 0} \frac{\gamma(\tilde{\eta})}{1-F(s(\tilde{\eta})|x)} = \infty.$$

Hence, there exists $\tilde{\eta} \in (0, \eta)$ such that $\gamma(\tilde{\eta})F(s|x) > 1-F(s(\tilde{\eta})|x)$. Fix any such $\tilde{\eta}$, and let $s = s(\tilde{\eta})$.

Given this value of $\tilde{\eta}$ and s (which, note, does not depend on δ), let

$$\beta = \frac{1}{\delta} \left(\frac{\gamma(\tilde{\eta})F(s|x)}{1-F(s|x)} - 1 \right) > 0.$$

Finally, let

$$\delta_1 = 1 - \frac{1-F(s|x)}{\gamma(\tilde{\eta})F(s|x)} < 1,$$

and note that

$$\frac{1-\delta_1}{\delta_1} = \frac{1-F(s|x)}{\gamma(\tilde{\eta})F(s|x) - (1-F(s|x))}.$$

Note that, for any $\delta \geq \delta_1$, we have

$$\alpha = 1 - (1 - \delta)\beta = 1 - \frac{1 - \delta}{\delta} \left(\frac{\gamma(\tilde{\eta}) F(s|x)}{1 - F(s|x)} - 1 \right) \geq 1 - \frac{1 - \delta_1}{\delta_1} \left(\frac{\gamma(\tilde{\eta}) F(s|x)}{1 - F(s|x)} - 1 \right) = 0,$$

so the strategy profile is well-defined.

By construction, we have

$$\begin{aligned} & F(s|x) + xF_x(s|x) + \beta\delta F_x(s|x)V \\ = & F(s|x) + xF_x(s|x) + \frac{\beta\delta x F_x(s|x) F(s|x)}{1 + \beta\delta(1 - F(s|x))} \\ = & F(s|x) - \frac{1 + \gamma(\tilde{\eta})}{\gamma(\tilde{\eta})} F(s|x)(1 - F(s|x)) + \frac{\left(\frac{\gamma(\tilde{\eta})F(s|x)}{1 - F(s|x)} - 1 \right) \left(-\frac{1 + \gamma(\tilde{\eta})}{\gamma(\tilde{\eta})} F(s|x)(1 - F(s|x)) \right) F(s|x)}{1 + \left(\frac{\gamma(\tilde{\eta})F(s|x)}{1 - F(s|x)} - 1 \right) (1 - F(s|x))} \\ = & 0. \end{aligned}$$

So the first-order condition of the claimant's problem holds, and (as we have seen) the second-order condition holds by log-concavity of F . Moreover,

$$V = \frac{x F(s|x)}{1 + \beta\delta(1 - F(s|x))} = \frac{x F(s|x)}{1 + \gamma(\tilde{\eta}) F(s|x) - (1 - F(s|x))} = \frac{x}{1 + \gamma(\tilde{\eta})} > \frac{x}{1 + \tilde{\eta}} > \frac{x}{1 + \eta} > x - \eta.$$

Similarly, the responder's value equals

$$\frac{(1 - x) F(s|x)}{1 + \beta\delta(1 - F(s|x))} = \frac{1 - x}{1 + \gamma(\tilde{\eta})} > 1 - x - \eta.$$

Finally, the responder's strategy is optimal so long as the one-shot gain from accepting the claimant's demand of x is outweighed by the lost continuation value of $\delta(1 - x)/(1 + \gamma(\tilde{\eta}))$: that is, if

$$(1 - \delta)(1 - x) \leq \delta \frac{1 - x}{1 + \gamma(\tilde{\eta})}, \text{ or } \delta \geq \delta_2 := \frac{1 + \gamma(\tilde{\eta})}{2 + \gamma(\tilde{\eta})}.$$

Taking $\bar{\delta} = \max\{\delta_1, \delta_2\}$ completes the proof.

A.7 Proof of Theorem 3

Fix any PPE and any $\eta > 0$. For any period t and public history h^t , let $U^t(h^t)$ denote the responder's continuation payoff at history h^t , and let $\bar{U} = \sup_{t,h^t} U^t(h^t)$. We show that $\bar{U} \leq \eta$. Since this holds for any $\eta > 0$, we have $\bar{U} \leq 0$.

Fix (t, h^t) such that $U^t(h^t) \geq \bar{U} - (1 - \delta)\eta$. Note that the responder's continuation payoff starting in period $t + 1$ is at most \bar{U} . Hence, if the responder accepts with probability 0 in period t conditional on reaching history h^t , then we have

$$U^t(h^t) \leq \delta \bar{U} \leq \delta (U^t(h^t) + (1 - \delta)\eta), \quad (8)$$

or equivalently $U^t(h^t) \leq \delta\eta$, which implies that $\bar{U} \leq \eta$. So suppose that the responder accepts with positive probability in period t conditional on reaching history h^t .

I claim that, in this case, the claimant uses a pure strategy at history h^t . Since the equilibrium strategies are public, each player's continuation payoff at the start of period $t + 1$ depends only on h^t , a_t (the current-period {accept, reject} decision), and z_t (which realizes after a_t). Hence, by the same argument as in the proof of Lemma 2, the responder uses a cutoff strategy in period t , with $s^* > \underline{s}$. Therefore, the claimant's expected utility from demanding x in period t equals $((1 - \delta)x + \delta V^A) F(s^*|x) + \delta V^R (1 - F(s^*|x))$, where V^A and V^R are her continuation payoffs from accept and reject, and s^* is the responder's cutoff signal. For any values of V^A and V^R , the same argument as in the proof of Lemma 2 shows that this expectation is unimodal in x . So the claimant uses a pure strategy.

Next, the claimant's demand at history h^t must leave the responder indifferent between accepting and rejecting. For, if the responder strictly preferred accepting, the claimant would increase her demand, which would increase her current-period payoff without affecting her continuation payoff (again because the equilibrium strategies are public). Thus, the responder's payoff is the same as that when he always rejects in period t . Hence, equation (8) applies, and we can again conclude that $\bar{U} \leq \eta$.

We have shown that $\bar{U} \leq 0$. Since the responder's minmax payoff is 0, this implies that, conditional on reaching any public history h^t , the responder's expected payoff in period t equals 0. Moreover, we have seen that, if the responder accepts with positive probability

conditional on reaching history h^t , the claimant must use a pure strategy that leaves the responder indifferent. Since the responder's continuation payoff at each history equals 0, this pure strategy must demand $x = 1$. By the full support assumption, it follows that either the claimant demands $x = 1$ with probability 1 or the responder rejects after every signal with probability 1.

A.8 Proof of Proposition 6

Suppose otherwise. Then for any $\varepsilon > 0$, there exists a period $T(\varepsilon)$ such that, for any $t \geq T(\varepsilon)$, the equilibrium probability that the responder accepts in period t (conditional on having accepted in all earlier periods) exceeds $1 - \varepsilon$. By the full support assumption, for any $\eta > 0$, there exists $\varepsilon(\eta) > 0$ such that if the responder's stage-game strategy $\sigma : S \rightarrow \Delta\{\text{accept, reject}\}$ satisfies $\int_S \sigma(s) f(s|x) ds > 1 - \varepsilon(\eta)$ for some $x \in [-M, M]$, then $\int_S \sigma(s) f(s|x = M) ds > 1 - \eta$. Hence, for any $\eta > 0$, if the claimant deviates to demanding $x_t = M$ in every period $t \geq T(\varepsilon(\eta))$, her continuation payoff starting from period $T(\varepsilon(\eta))$ is at least

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t (1 - \eta)^{t+1} M = \frac{(1 - \delta)(1 - \eta)M}{1 - \delta(1 - \eta)}.$$

This expression is strictly greater than 1 iff

$$\eta < \eta^* := \frac{(1 - \delta)(M - 1)}{M - \delta(M - 1)},$$

which is strictly positive because $\delta \in (0, 1)$ and $M > 1$. Therefore, for any $\eta < \eta^*$, the claimant can guarantee a continuation payoff strictly greater than 1 starting from period $T(\varepsilon(\eta))$. However, the sum of the players' payoffs in each period is at most 1, and the responder's equilibrium continuation payoff starting from any period is at least her minmax payoff of 0. So this is a contradiction.

A.9 Proof of Theorem 4

We first note that, by the same argument as in Lemma 1, a player's demand cannot be accepted with equilibrium probability 1. Hence, each player must be indifferent between

accepting and rejecting her opponent's demand in any SCE, so every SCE must satisfy the Rubinstein equations, (3). By the argument given in the text, it follows that, in every SCE, (p_i, p_j) must satisfy (4) (as well as the symmetric equation for player j).

Since $g_i(z)/G_i(z)$ is decreasing in z (as $F_{i,x}(s_i|x_i)/F_i(s_i|x_i) = -g_i\left(\frac{s_i-x_i}{\theta_i}\right)/G_i\left(\frac{s_i-x_i}{\theta_i}\right)$ is increasing in s_i , by Lemma 4), letting $z = G_i^{-1}(1-p_i)$, we see that $g_i(G_i^{-1}(1-p_i))/(1-p_i)$ is increasing in p_i , and therefore the left-hand side of (4) is increasing in p_i , as well as increasing in p_j . Since the left-hand side of (4) is clearly continuous in p_i and p_j , this equation defines a continuous, downward-sloping curve C_i in (p_i, p_j) -space. Let \underline{p}_i and \bar{p}_i be the smallest and greatest values of \bar{p}_i lying on this curve, which are given by

$$\frac{1 - \delta^2}{1 - \delta^2 \underline{p}_i} \frac{g_i\left(G_i^{-1}\left(1 - \underline{p}_i\right)\right)}{1 - \underline{p}_i} = \theta_i \text{ and } (1 - \delta) \frac{g_i\left(G_i^{-1}\left(1 - \bar{p}_i\right)\right)}{1 - \bar{p}_i} = \theta_i.$$

(That is, the points $(\underline{p}_i, 1)$ and $(\bar{p}_i, 0)$ lie on the curve C_i .) Note that $\underline{p}_i > 0$ and $\bar{p}_i < 1$ because, by the informativeness assumption,

$$\begin{aligned} \lim_{p_i \rightarrow 0} \frac{1 - \delta^2}{1 - \delta^2 p_i} \frac{g_i\left(G_i^{-1}\left(1 - p_i\right)\right)}{1 - p_i} &= \lim_{z \rightarrow \bar{z}} (1 - \delta^2) g_i(z) < \theta_i \text{ and} \\ \lim_{p_i \rightarrow 1} (1 - \delta) \frac{g_i\left(G_i^{-1}\left(1 - p_i\right)\right)}{1 - p_i} &= \lim_{z \rightarrow \underline{z}} (1 - \delta) \frac{g_i(z)}{G_i(z)} > \theta_i. \end{aligned}$$

Define the curve C_j and $0 < \underline{p}_j < \bar{p}_j < 1$ symmetrically: in particular, the points $(0, \bar{p}_j)$ and $(1, \underline{p}_j)$ lie on the curve C_j . Since $0 < \underline{p}_i < \bar{p}_i < 1$, $0 < \underline{p}_j < \bar{p}_j < 1$, and the curves C_i and C_j are continuous, by the intermediate value theorem they intersect at some point (p_i, p_j) satisfying $\underline{p}_i < p_i < \bar{p}_i$ and $\underline{p}_j < p_j < \bar{p}_j$. Moreover, any point (p_1, p_2) at the intersection of C_i and C_j , together with the corresponding demands (x_1, x_2) given by the Rubinstein equations, (3), corresponds to an SCE, because each player is indifferent between accepting and rejecting the opponent's demand, and each player's demand is optimal given the opponent's acceptance rule. (Conversely, any tuple (p_1, p_2, x_1, x_2) that does not satisfy these conditions is not an SCE.) This proves part 1 of the theorem.

Since the curves C_i and C_j are both downward-sloping, for any two intersections (p_1, p_2) and (p'_1, p'_2) with $p'_1 \geq p_1$, we have $p'_2 \leq p_2$. Letting (x_1, x_2) and (x'_1, x'_2) be the corresponding

solutions to the Rubinstein equations, we have $x'_1 \leq x_1$ and $x'_2 \geq x_2$. By the formulas for the players' equilibrium payoffs, (2), the facts that $p'_1 \geq p_1$, $x'_1 \leq x_1$, $x'_2 \geq x_2$, $x_1 \geq \delta(1 - x_2)$, and $x'_1 \geq \delta(1 - x'_2)$ (where the latter two inequalities were established in the text) imply that $U'_1 \leq U_1$ and $U'_2 \geq U_2$. This proves part 2 of the theorem.

Finally, let (p_1^*, p_2^*) and (p_1^{**}, p_2^{**}) be the extremal SCE values of (p_1, p_2) . Fix $i \in \{1, 2\}$. Slightly abusing notation, for each $p_i \in [\underline{p}_i, \bar{p}_i]$, let $C_i(p_i)$ (resp., $C_j(p_i)$) denote the value of p_j such that (p_i, p_j) lies on the curve C_i (resp., C_j). Note that $C_i(\underline{p}_i) = 1 > C_j(\underline{p}_i)$ and $C_i(\bar{p}_i) = 0 < C_j(\bar{p}_i)$. Note also that increasing θ_i shifts the curve C_i upwards and leaves the curve C_j fixed. Applying Theorem 1 of Milgrom and Roberts (1994) to the function $C_i(p_i) - C_j(p_i)$, it follows that increasing θ_i results in an increase in both p_1^* and p_1^{**} , and hence (since C_j is fixed and downward-sloping) a decrease in both p_2^* and p_2^{**} . Together with the Rubinstein equations, (3), and the formulas for the players' equilibrium payoffs, (2), both of which are independent of θ_i , it follows that the smallest and largest SCE values of x_i and U_i (resp., x_j and U_j) are decreasing (resp., increasing) in θ_i . This proves part 3 of the theorem.

References

- [1] Abreu, Dilip, and David Pearce. "A Dynamic Reinterpretation of Nash Bargaining with Endogenous Threats." *Econometrica* 83.4 (2015): 1641-1655.
- [2] Allison, Roy. *Russia, the West, and Military Intervention*. Oxford University Press, Oxford, 2013.
- [3] Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere. "Bargaining with Incomplete Information." *Handbook of Game Theory with Economic Applications* 3 (2002): 1897-1945.
- [4] Baliga, Sandeep, Ethan Bueno De Mesquita, and Alexander Wolitzky. "Deterrence with Imperfect Attribution." *American Political Science Review* 114.4 (2020): 1155-1178.
- [5] Bagwell, Kyle. "Commitment and Observability in Games." *Games and Economic Behavior* 8.2 (1995): 271-280.
- [6] Bhaskar, Venkataraman "Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems." *Review of Economic Studies* 65.1 (1998): 135-149.

- [7] Bhaskar, Venkataraman, and Ichiro Obara. "Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring." *Journal of Economic Theory* 102.1 (2002): 40-69.
- [8] Binmore, K. G.: "Perfect Equilibria in Bargaining Models," in *The Economics of Bargaining*, ed. by K. G. Binmore and P. Dasgupta. *Oxford: Basil Blackwell* (1987): 77-106.
- [9] Bolton, Patrick, and Antoine Faure-Grimaud. "Satisficing Contracts." *Review of Economic Studies* 77.3 (2010): 937-971.
- [10] Carlsson, Hans. "A Bargaining Model where Parties Make Errors." *Econometrica*, 59.5 (1991): 1487-1496.
- [11] Coy, Peter. "What Game Theory Says about China's Strategy," *New York Times*, Sept. 13, 2021.
- [12] Denti, Tommaso, Massimo Marinacci, and Aldo Rustichini. "Experimental Cost of Information." *American Economic Review* Forthcoming (2022).
- [13] Farrell, Joseph, and Robert Gibbons. "Cheap Talk Can Matter in Bargaining." *Journal of Economic Theory* 48.1 (1989): 221-237.
- [14] Fearon, James D. "Domestic Political Audiences and the Escalation of International Disputes." *American Political Science Review* (1994): 577-592.
- [15] Fearon, James D. "Rationalist Explanations for War." *International Organization* 49 (1995): 379-379.
- [16] Ferejohn, John. "Incumbent Performance and Electoral Control." *Public Choice* 50.1 (1986): 5-25.
- [17] Freedman, Lawrence. *Ukraine and the Art of Strategy*. Oxford University Press, 2019.
- [18] Fudenberg, Drew, David Levine, and Jean Tirole, "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information." *Game-Theoretic Models of Bargaining*, Cambridge University Press, West Nyack, NY (1985): 73-98.
- [19] Fudenberg, Drew, David Levine, and Eric Maskin. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62.5 (1994): 997-1039.
- [20] Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. "Foundations of Dynamic Monopoly and the Coase Conjecture." *Journal of Economic Theory* 39.1 (1986): 155-190.
- [21] Hart, Oliver D., and Jean Tirole. "Contract Renegotiation and Coasian Dynamics." *Review of Economic Studies* 55.4 (1988): 509-540.
- [22] Holmström, Bengt. "Moral Hazard and Observability." *Bell Journal of Economics* 10.1 (1979): 74-91.

- [23] Johnson, Justin P., and David P. Myatt. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review* 96.3 (2006): 756-784.
- [24] Kaplan, Robert D. *Asia's Cauldron: The South China Sea and the End of a Stable Pacific*. Random House, 2014.
- [25] Klein, Benjamin, and Keith B. Leffler. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy* 89.4 (1981): 615-641.
- [26] Maggi, Giovanni. "The Value of Commitment with Imperfect Observability and Private Information." *RAND Journal of Economics* (1999): 555-574.
- [27] Matsushima, Hitoshi. "On the Theory of Repeated Games with Private Information: Part I: Anti-Folk Theorem without Communication." *Economics Letters* 35.3 (1991): 253-256.
- [28] Merlo, Antonio, and Charles Wilson. "A Stochastic Model of Sequential Bargaining with Complete Information." *Econometrica* 63.2 (1995): 371-399.
- [29] Milgrom, Paul R. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* (1981): 380-391.
- [30] Milgrom, Paul, and John Roberts. "Price and Advertising Signals of Product Quality." *Journal of Political Economy* 94.4 (1986): 796-821.
- [31] Milgrom, Paul, and John Roberts. "Comparing Equilibria." *American Economic Review* (1994): 441-459.
- [32] Nash, John. "Two-Person Cooperative Games." *Econometrica* 21.1 (1953): 128-140.
- [33] Powell, Robert. "Uncertainty, Shifting Power, and Appeasement." *American Political Science Review* (1996a): 749-764.
- [34] Powell, Robert. "Bargaining in the Shadow of Power." *Games and Economic Behavior* 15.2 (1996b): 255-289.
- [35] Prat, Andrea, and David Strömberg. "The Political Economy of Mass Media." *Advances in Economics and Econometrics* 2 (2013): 135.
- [36] Ravid, Doron. "Ultimatum Bargaining with Rational Inattention." *American Economic Review* 110.9 (2020): 2948-63.
- [37] Rey, Patrick, and Jean Tirole. "A Primer on Foreclosure." *Handbook of Industrial Organization* 3 (2007): 2145-2220.
- [38] Rubinstein, Ariel. "Perfect Equilibrium in a Bargaining Model." *Econometrica*, 50.1 (1982): 97-109.
- [39] Sanger, David E. *The Perfect Weapon: War, Sabotage, and Fear in the Cyber Age*. Crown, New York, 2019.

- [40] Schelling, Thomas C. *The Strategy of Conflict*. Harvard University Press, Cambridge, 1960.
- [41] Schelling, Thomas C. *Arms and influence*. Yale University Press, New Haven, 1966.
- [42] Sekiguchi, Tadashi. "Efficiency in Repeated Prisoner's Dilemma with Private Monitoring." *Journal of Economic Theory* 76.2 (1997): 345-361.
- [43] Shapiro, Carl, and Joseph E. Stiglitz. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74.3 (1984): 433-444.
- [44] Shortle, James S., and Richard D. Horan. "The Economics of Nonpoint Pollution Control." *Journal of Economic Surveys* 15.3 (2001): 255-289.
- [45] Tirole, Jean. "Cognition and Incomplete Contracts." *American Economic Review* 99.1 (2009): 265-94.
- [46] Van Damme, Eric, and Sjaak Hurkens. "Games with Imperfectly Observable Commitment." *Games and Economic Behavior* 21.1-2 (1997): 282-308.
- [47] Wolinsky, Asher. "Prices as Signals of Product Quality." *Review of Economic Studies* 50.4 (1983): 647-658.