Noisy Screening and Brinkmanship*

Germán Gieczewski†

August 2023
PRELIMINARY AND INCOMPLETE.

Abstract

I study a repeated bilateral relationship subject to termination. One player, the proposer, offers a transfer in each period. The receiver can accept and continue the relationship, or quit and take an outside option whose value is private information. Unlike in Coasian bargaining, remaining types are those from whom more can be extracted, potentially inviting ratcheting. Tirole (2016) shows that, if the receiver's type is persistent, there is actually no ratcheting. I show that, if the receiver's incentives to accept are affected even by small, transient shocks that the proposer cannot perfectly observe, then offers may worsen over time, until the receiver inevitably quits. The reason is that a small escalation causes exit only if the receiver's type is marginal *and* the shock is unfavorable. Major escalations may alternate with periods of slow ratcheting. Exit may be inevitable even if the proposer has commitment power. Applications include crisis bargaining in international relations and surplus extraction from an employee.

*I am grateful to Arjada Bardhi, Dan Bernhardt, Nina Bobkova, Alessandra Casella, Micael Castanheira, Sylvain Chassang, Kirill Evdokimov, Andreas Kleiner, Stefan Krasa, Matías Iaryczower, Kristopher Ramsay, Peter Schram, Alastair Smith, Juuso Toikka, and participants at the 2023 Benelux PECO, SITE Political Economy Session, and MPSA and APSA Annual Meetings for helpful comments and suggestions.

†Department of Politics, Princeton University. Address: 311 Fisher Hall, Princeton NJ 08540, USA. Phone number: 6172330642. E-mail: `germang@princeton.edu`.

1 Introduction

This paper concerns itself with settings in which a proposer (he) makes offers over time to a receiver (she) with a persistent and privately known outside option. A crucial feature of the setting is that acceptance simply continues the game, with the current offer determining payoffs only for one period, whereas rejection ends the relationship permanently. In such settings, the proposer might infer from the acceptance of a stingy offer that the receiver can be further taken advantage of.

In principle, the proposer may then want to ratchet the aggressiveness of his offers over time. If this ratcheting is gradual, it may appear to take the form of "salami tactics", eloquently described by Schelling (1966, p. 66–67):

> "'Salami tactics', we can be sure, were invented by a child . . . Tell a child not to go in the water and he'll sit on the bank and submerge his bare feet; he is not yet 'in' the water. Acquiesce, and he'll stand up; no more of him is in the water than before. Think it over, and he'll start wading, not going any deeper; take a moment to decide whether this is different and he'll go a little deeper, arguing that since he goes back and forth it all averages out. Pretty soon we are calling to him not to swim out of sight, wondering whatever happened to all our discipline."

Tirole (2016) shows that this intuition is incorrect: in fact, even when the proposer can update his offer over time, it is optimal to make the same offer as would be made in a one-shot version of the game, and then repeat it in every period. The receiver, too, behaves "myopically", accepting or rejecting based on flow payoffs, and there is thus no screening beyond the first period: all receiver types who quit do so in the first period. The reason is that, when the proposer considers making a slightly more aggressive offer than the initial one, the marginal trade-off that arises—most of the time, the proposer extracts a little more, but a few more marginal types will reject—is the same regardless of whether low, infra-marginal types have already screened out or not. In addition, the proposer's commitment solution is time-consistent, and hence it is also the proposer-optimal equilibrium of the game with no commitment. These results hold even if the players' discount factors differ, so long as the receiver is not strictly more patient than the proposer.

The general thrust of this paper is to show that these results may change dramatically when the screening process is beset even by small amounts of noise. In the simplest perturbation of the model, either the receiver's payoff from accepting the proposal, or the value of her outside option, is affected by a small, transient shock that is re-drawn in every period, and which cannot be seen by the proposer before he chooses his offer. In this model, I first show that if the receiver is completely myopic, then the proposer's optimal path of offers is the same whether he has commitment power or not—but, rather than being constant over time, it features increasingly aggressive offers which eventually induce every type of the receiver to quit.[1]

The logic behind this result is as follows. In the presence of noise, ensuring that no remaining receiver type will quit requires the proposer to make her offer more generous, by just enough to offset the worst possible shock realization. Making her offer slightly *less* generous than that will be punished by the receiver exiting only when two unlikely events coincide: the receiver's type is marginal *and* the shock realization is very unfavorable. For a small enough escalation by the proposer, the risk he runs is thus second-order, and dominated by the direct gain of extracting a little more from the receiver. Once some marginal types have been "cut" from the distribution, the proposer is tempted to further increase her demands, continuing the process until almost all receiver types have quit.

The proposer's optimal offer "slices" off only a small fraction of marginal receiver types whenever the optimal offer in the world without shocks would induce no screening at all. But progressively truncating the bottom of the distribution in this way can eventually induce a distribution of posteriors for which the optimal offer, with or without shocks, must screen a large set of types. The path of offers can thus alternate between long periods of slow ratcheting and discrete jumps that induce a high immediate risk of exit by the receiver.

Next, I consider the more general case in which the receiver is also forward-looking. In general, the commitment solution is no longer time-consistent. However, it, too, leads to eventual exit of all receiver types, so long as the receiver is *strictly* more impatient than the proposer. The logic is that, when the proposer hypothetically increases her demand in a period t, there are two costs: a higher probability that the receiver will quit in period t, and a higher probability that the receiver will quit in periods *before* t, due to expecting a lower continuation value from remaining in the

[1]Notably, Tirole (2016) shows that shocks observed by the proposer before offers are made do not alter his results substantially. Hence, what is crucial is that unobserved shocks make the receiver's behavior slightly less predictable from the proposer's point of view.

game, in particular in period t. A proposer choosing a path of offers *ex ante* with commitment power must weigh both costs. It turns out that the former cost, as in the case of a myopic receiver, cannot deter the proposer. The latter cost can, and does—but only if the receiver is as patient as or more patient than the proposer. If the receiver is more impatient than the proposer, she cares less about escalations made in much later periods, and hence reacts little to them in advance, so that these "proactive punishments" are not a concern for the proposer.

Finally, I study the problem when the receiver is forward-looking and the proposer does not have commitment power. Here, although the proposer's high demands in a period t may retroactively hurt him in earlier periods, he cannot commit not to indulge in such escalations ex post. According to this intuition, salami-slicing ought to be more likely in the absence of commitment power. However, a new force can now dissuade the proposer: the receiver may believe after a small deviation today that, in the continuation, the proposer will make much more aggressive demands than he would have made on the equilibrium path. This belief may lead to increased exit today, punishing deviations by the proposer. I show that this force may prevent full screening, but its bite disappears if either the receiver cannot perfectly observe deviations by the proposer, or if we restrict attention to equilibria that are Markovian in a certain sense.

Our model captures the features of several salient examples of repeated bargaining in the shadow of conflict, including ones in which salami tactics have clearly been employed. For instance, China has attempted to gradually expand its control over the South China Sea by building artificial islands, and blocking access to existing islands by surrounding them with boats (Himmelman, 2013). Similar tactics have been used by Israeli settlers gradually encroaching on the West Bank (Krause and Eiran, 2018). Conversely, China has accused the United States of gradually advancing ties with Taiwan (Coy, 2021), while Russia has claimed that it was forced to invade Ukraine by a gradual encroachment by the West on Ukraine (Mearsheimer, 2022). The logic of our model is also present in employment relationships in which an employer tries to extract more and more from a worker over time, by lowering pay or demanding more output for the same work, and in which the worker's outside option is unknown by the employer.[2] Much the same problem is faced by the perpetrators of a protection

[2] In the closely related "ratchet effect" literature (Freixas, Guesnerie and Tirole, 1985; Laffont and Tirole, 1988; Hart and Tirole, 1988), the worker chooses how hard to work, and may under-exert

racket who can update their protection fees over time, or by Muslim governments that historically levied the *jizya* tax on non-Muslims, who could avoid it by permanently converting to Islam (Tirole, 2016; Saleh, 2018; Saleh and Tirole, 2021).

This paper is connected mainly to two literatures. First, in Coasian bargaining models (Gul, Sonnenschein and Wilson, 1986), a monopolist chooses prices without commitment and faces a buyer whose valuation is unknown. If high-valuation buyers screen out by buying at a high price, the monopolist is then tempted to lower her price so as to sell to low types—but then even high types, anticipating this, reject low offers. The classic Coase conjecture is that, as a result, the monopolist sells to all types at the lowest valuation in the support. This conjecture holds in the benchmark setting (Gul et al., 1986), but may fail to hold if the players have interdependent values (Deneckere and Liang, 2006; Fuchs and Skrzypacz, 2013b); if the seller's cost varies stochastically and the set of buyer types is discrete (Ortner, 2017); if traders or information arrives over time (Fuchs and Skrzypacz, 2010); if the seller has private information about her beliefs about the buyer's valuation (Feinberg and Skrzypacz, 2005); or if the players have different terminal payoffs, either as the result of deadlines (Fuchs and Skrzypacz, 2013a), access to outside options (Board and Pycia, 2014; Hwang and Li, 2017), or the possibility that either player may "collapse", as may happen if the bargaining process models a war (Baliga and Sjöström, 2023).

In most models of Coasian bargaining there is *negative* selection: remaining types are those from whom less can be extracted.[3] In contrast, in our model, remaining types are those who the proposer can exploit, so there is *positive* selection, as in Tirole (2016) and Saleh (2018). A few recent papers (Ali, Kartik and Kleiner, n.d.; Evdokimov, 2023) instead consider sequential bargaining with incomplete information in which acceptance ends the game but the receiver has single-peaked preferences, as in the canonical agenda-setting game (Romer and Rosenthal, 1978, 1979). Besides considering screening problems of a different nature, the extant literature does not consider the effects of transient shocks to the receiver's preferences, as I do.

Second, in crisis bargaining models (Fearon, 1995), a proposer offers to split a resource in some fashion, and a receiver can accept the offer or go to war, leading to some exogenous payoffs about which she is privately informed. Models of screening

effort to avoid revealing a high type from which more can be extracted. In contrast, in our setting, the employer sets the bar and the worker has only two options: meet the demands or quit.

[3]A partial exception is models in which the buyer can quit either by accepting or taking an outside option (Board and Pycia, 2014; Hwang and Li, 2017).

with positive selection, such as Tirole (2016) and this paper, can be seen as models of repeated crisis bargaining. The existing literature considers mostly static generalizations of the benchmark model (e.g., Fey and Ramsay 2011). Fey, Meirowitz and Ramsay (2013) shows a similar result to Tirole (2016)'s in a two-period model, while Fey and Kenkel (2021) consider a setting in which the receiver can accept an offer *permanently*; fight (also leading to a permanent outcome); or reject and continue negotiations, leading to alternating offers.

A strand of the crisis bargaining literature considers transfers that may themselves affect the balance of power, i.e., the receiver's outside option (Fearon, 1996; Powell, 2006; Schwarz and Sonin, 2008). In these models, a different form of "salami tactics" may arise: the receiver may acquiesce to a path of front-loaded offers that leaves her progressively worse off, because early payoffs are high enough to compensate for low payoffs later—and, by the time payoffs decline, fighting has become too costly (i.e., the outside option has deterioriated too much). In our model, in contrast, the proposer becomes more aggressive over time even though the receiver is not weakened—and eventually chooses to fight.

Finally, our model is also related to the aforementioned literature on ratchet effects, as in both types of models the proposer would like to take advantage of a high type. The difference is that in the ratcheting literature, the agent cannot quit permanently, so there is no selection; instead, good types may collect rents by pooling with bad types. Under natural assumptions, this may lead to the principal not being able to ever learn the agent's type (Laffont and Tirole, 1988; Hart and Tirole, 1988). Acharya and Ortner (2017) show that the agent may reveal her type in a ratcheting model if the environment changes over time as a result of shocks. Their model assumes commonly observed shocks (which would not produce much revelation in our setting, as shown by Tirole (2016))—and large enough that production may be inefficient. Moreover, their result is driven by a different logic: essentially, good types can only extract rents in good states, so they may be willing to reveal their type when future states are unlikely to be good.

2 The Model

Time is discrete: $t = 0, 1, \dots, T$, where T may be finite or infinite. There are two players: 1, the *proposer*, and 2, the *receiver*, with discount factors $\delta_1, \delta_2 \in [0, 1)$. In each period t, the proposer makes a *demand* $x_t \in \mathbb{R}$. The receiver can accept it or quit. If the receiver accepts, flow payoffs are generated as a function of x_t and the game continues on to the next period. If the receiver quits, she receives an outside option whose value depends on a parameter θ. θ is drawn from a c.d.f. F admitting a continuous density f with support $[\underline{\theta}, \overline{\theta}]$, where $0 < \underline{\theta} < \overline{\theta}$. It is drawn and shown to the receiver at the beginning of the game.

The Unperturbed Model

Before presenting the full model, it is useful to describe the *unperturbed* model that has been studied in the literature. In the unperturbed model, which is formally equivalent to the benchmark model in Tirole (2016) or to a repeated crisis bargaining model, flow payoffs from a proposal x_t are of the form $(\pi(x_t), -x_t)$, where $\pi : [0, +\infty) \to \mathbb{R}$ is a real-valued, smooth, strictly increasing and weakly concave function, with $\pi(0) \geq 0$. Rejection ends the interaction, locking in flow payoffs $(0, -\theta)$ in the current period as well as any remaining periods. The assumptions $\underline{\theta} > 0$, $\pi(0) \geq 0$ imply that, regardless of the receiver's type, exit is inefficient, as both players would do better if an offer $x_t = 0$ were made and accepted in every period.

In the case of an employment relationship, x_t represents the amount of effort demanded by the employer, or a bundle of effort and diminished pay. θ is the worker's flow loss from taking the outside option, relative to the current relationship (under the ideal offer path $x_t \equiv 0$). Higher types are thus more willing to accept offers and more exploitable by the proposer. If we add θ to the receiver's flow payoff, her flow payoff from acceptance is $\theta - x_t$ and her payoff from rejection is a fixed number, as in the benchmark model in Tirole (2016). The payoffs are also equivalent to those in a crisis bargaining model up to an affine transformation, where x_t represents the share of the resource demanded by player 1; $\pi(x_t)$ is player 1's valuation of this share; $1 - x_t$ is the share offered to player 2; θ is the receiver's cost of fighting a war with a permanent outcome.

In the one-shot version of the game ($T = 0$), the receiver accepts a proposal x iff $\theta \geq x$. The proposer's payoff from an offer x is then $u(x) = \pi(x)(1 - F(x))$. Let x^* be a maximizer of u. If x^* is an interior optimum, it must satisfy the FOC

$$\pi'(x)(1 - F(x)) = \pi(x)f(x), \tag{1}$$

where $\pi'(x)$ is the marginal benefit of an increased demand, extracted from all the complying types, who make up a mass of size $1 - F(x)$, while $\pi(x)$ is the amount lost for every marginal type who now rejects, and $f(x)$ is the density of marginal types.

It then follows from Propositions 1-3 in Tirole (2016) that, if x^* is unique and $\delta_2 \leq \delta_1$, then, in the full commitment solution, the proposer sets $x_t \equiv x^*$ for all t, and the receiver accepts at all times if $\theta \geq x^*$ and rejects right away otherwise. Moreover, this outcome is time-consistent, so it can also be obtained in a Perfect Bayesian equilibrium. When $\delta_2 > \delta_1$, this strategy profile is still a PBE, although it is no longer the full commitment solution (as the proposer may prefer to commit to backloading incentives, making very low demands in late periods which will keep the receiver in the game through early periods even with high demands.)

Proposition 1 (Tirole 2016 Propositions 1+2+3)**.** *Suppose $T < \infty$ and x^* is unique. For any δ_1, δ_2, in the unique PBE, the proposer sets $x_t = x^*$ for all t.*

If $\delta_1 \geq \delta_2$, then the proposer can do no better with commitment power. (If $\delta_1 < \delta_2$, the proposer can do better under commitment by backloading payoffs.)

The Full Model

We now modify the preceding model by assuming that, in each period t, the receiver's cost of agreeing to a demand x_t is affected by a shock ϵ_t, so that flow payoffs in period t from acceptance are $(\pi(x_t), -x_t - \epsilon_t)$. Payoffs from rejection are unchanged. ϵ_t is drawn by Nature from a c.d.f. G and shown to the receiver before her decision, but never observed by the proposer. (Later on, we relax this assumption to allow for the proposer to learn about the shock *after the fact*, i.e., at the end of the period). Shocks are drawn independently across periods and independently of θ. Throughout the paper, we will assume that, for some value of a parameter $\eta > 0$, G satisfies one of two possible assumptions:

A1(η) G admits a density g which is symmetric around 0, has support $[-\eta, \eta]$, and is continuous in $[-\eta, \eta]$.

A2(η) G satisfies A1(η) and, in addition, $g(\eta) = 0$.

Our assumptions about F and G rule out probability masses in either distribution. If either distribution features probability masses, the proposer's problem becomes much closer to what it would be either if he knew the receiver's type or if there were no shocks to the receiver's flow payoffs. As a running example of a distribution satisfying A1(η), take $\epsilon_t \sim U[-\eta, \eta]$.

A distribution satisfying A2(η) may be a "smoothed out" perturbation of any density satisfying A1(η). Alternatively, if the shock ϵ_t is in fact the sum of any $k > 1$ shocks ϵ_{it} that are independently drawn from distributions satisfying A1(η_i) with $\eta_i > 0$ and $\sum_i \eta_i = \eta$, then ϵ_t would satisfy A2(η). For instance, if ϵ_t were the sum of two uniform shocks with support $\left[-\frac{\eta}{2}, \frac{\eta}{2}\right]$, its c.d.f. would be triangular with support $[-\eta, \eta]$, thus satisfying A2(η). The additional assumption $g(\eta) = 0$ imposed by A2(η) ensures that, from the proposer's point of view, a draw of ϵ_t near the top of the distribution is very unlikely.

Additionally, it is worth noting that shocks to the receiver's acceptance payoff are equivalent to shocks to her outside option, in the following sense. Suppose we assumed that flow payoffs from acceptance at time t were $(\pi(x_t), -x_t)$ as in the unperturbed model, but that the receiver's outside option at time t generated a flow payoff $-\theta + \epsilon_t$ in period t, and the usual θ_t thereafter. To see the logic clearly, suppose T is finite. For $t = T$, the receiver's incentives are as in our main specification, as her payoffs are the same, plus an exogenous shock ϵ_T which she receives irrespective of her action. At time $T - 1$, the same logic applies, except that her continuation value from accepting is now affected by the shock ϵ_T which she will receive at time T. However, since $E(\epsilon_T) = 0$ by assumption, this shock is zero in expectation. Thus her incentives are the same at time $T - 1$, and so on.

More generally, it follows that the model and its solution are fundamentally unchanged if we assume that the outside option is subject to shocks instead of the receiver's flow payoff from acceptance. In a principal-agent relationship, a shock to the outside option reflects that the worker receives slightly better or worse offers in different periods, while a shock to her acceptance payoff may represent that the task demanded by the employer is slightly harder or easier to complete than the employer anticipated. In a crisis bargaining context, the natural interpretation is that a shock to the outside option represents changes in the receiver's cost of fighting, while a shock to acceptance payoffs represents changes either in the receiver's valuation, e.g., of territory in dispute, or in the audience costs of accepting a given demand.

Finally, in the perturbed version of the model, we will strengthen the assumption $\underline{\theta} > 0$ by assuming further that $\underline{\theta} > \eta$, if G satisfies A$i(\eta)$. This assumption ensures that the proposer can always ensure acceptance by proposing $x_t = 0$, regardless of the realizaton of ϵ_t.

We will consider two versions of the problem. In the *full commitment* version, the proposer credibly commits to a sequence of demands $(x_t)_{t\geq 0}$ at the beginning of the game. The receiver then faces a decision theory problem. In the *no commitment version*, the proposer chooses a demand x_t at the beginning of period t and cannot make any commitments regarding future demands.

Our solution concept is Perfect Bayesian equilibrium.

3 Myopic Receiver

We begin by fully characterizing the solution in the special case where the receiver is completely myopic, that is, $\delta_2 = 0$. As we will see, in this case, the proposer's optimal demand path in the full commitment problem is time-consistent, and hence the proposer attains the same outcome in the game without commitment. However, in contrast to Tirole (2016)'s setting, in which the optimal demand path is constant, here the proposer escalates until he forces the receiver to quit.

Proposition 2. *Suppose $\delta_2 = 0$, $T = \infty$, and G satisfies A1(η) for any $\eta > 0$. Then:*

(i) In all PBEs of the game without commitment, the proposer attains his full-commitment payoff. Moreover, a demand path $(x_t)_t$ is compatible with PBE iff it attains this payoff.

(ii) In any equilibrium, either with or without commitment, $\liminf_{t\to\infty} x_t \geq \overline{\theta} - \eta$. Hence, the probability that the receiver exits on the equilibrium path is 1.

The logic behind the proposition is instructive, as it underpins much of our later results. Since the receiver is myopic, she accepts a demand x_t at time t if and only if $\theta \geq x_t + \epsilon_t$. (It is without loss to assume that the receiver accepts when indifferent.) The proposer then effectively faces a decision theory problem, namely,

$$\max_x \sum_{t=0}^{\infty} \delta_1^t \pi(x_t) \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_t(\theta; x) d\theta,$$

where $x = (x_t)_t$, and $P_t(\theta; x)$ is the probability that a receiver of type θ accepts all demands *through* period t inclusive. Since the receiver accepts at time t iff $\theta \geq x_t + \epsilon_t$,

$$P_t(\theta; x) = \prod_{s=0}^{t} G(\theta - x_s).$$

Part (i) follows from the fact that the choice of $(x_t)_{t \geq s}$ has no impact on the receiver's behavior before period s.

Let $f_t(\theta) = f(\theta) P_{t-1}(\theta; x)$ be the density of receiver types left in the game at the beginning of period t, and $F_t(\theta)$ be the corresponding c.d.f. Let $U_t(\theta; x)$ be the proposer's expected continuation utility at the beginning of period t under a demand path x, conditional on the receiver's type being θ. The proposer's objective is thus $E(U_0(\theta; x))$, with the expectation taken over θ. The marginal impact of shifting x_t, conditional on a receiver type θ, is (proportional to)

$$\delta_1^{-t} \frac{\partial U_0(\theta; x)}{\partial x_t} = P_t(\theta; x) \pi'(x_t) - P_{t-1}(\theta; x) \left(\pi(x_t) + \delta_1 U_{t+1}(\theta; x) \right) g(\theta - x_t).$$

The impact unconditional on the receiver's type is thus (proportional to)

$$\delta_1^{-t} \frac{\partial E(U_0(\theta; x))}{\partial x_t} = \pi'(x_t) \int_{\underline{\theta}}^{\overline{\theta}} f_t(\theta; x) G(\theta - x_t) d\theta - \int_{\underline{\theta}}^{\overline{\theta}} f_t(\theta; x) \left(\pi(x_t) + \delta_1 U_{t+1}(\theta; x) \right) g(\theta - x_t) d\theta.$$

If x_t is an interior optimum, this yields an analogous condition to the one-shot FOC (1):

$$\pi'(x_t) \int_{\underline{\theta}}^{\overline{\theta}} f_t(\theta; x) G(\theta - x_t) d\theta = \int_{\underline{\theta}}^{\overline{\theta}} f_t(\theta; x) \left(\pi(x_t) + \delta_1 U_{t+1}(\theta; x) \right) g(\theta - x_t) d\theta. \quad (2)$$

Now part (ii) follows from the following argument. Note that, since $P_t(\theta)$ is weakly decreasing in t for each θ, and weakly increasing in θ, $P_t(\theta) \searrow P_\infty(\theta)$ for some function $P_\infty(\theta)$. Suppose that the receiver does not exit with probability 1. Equivalently, suppose that $P_\infty(\theta) > 0$ for some $\theta < \overline{\theta}$. Let $\theta_\infty = \inf(\text{supp}(P_\infty(\theta)))$. Then $\theta_\infty = \limsup_t x_t + \eta$.

For large t, the left-hand side of (2), which measures the marginal gain from increasing x_t, is positive. Indeed, it is bounded below by $\pi'(x_t) F_\infty(\overline{\theta})$, which is bounded away from zero, as optimal offers cannot grow without bound—in fact, they can never be greater than $\overline{\theta} + \eta$. On the other hand, the right-hand side, which measures the marginal cost stemming from additional rejections, must converge to zero, since $f_t(\theta; x)$ converges pointwise to zero below θ_∞, $\pi(x_t) + \delta_1 U_{t+1}(\theta; x)$ is bounded, and $g(\theta - x_t)$ goes to zero for all $\theta > \theta_\infty$, a contradiction.

A more informal intuition goes as follows: for t large enough that most receiver types who *would* have quit have already done so, the proposer can guarantee that virtually no more receiver types *will* quit if he proposes any $x_t \leq \theta_\infty - \eta$. However, if he pushes a little beyond that, demanding $x_t = \theta_\infty - \eta + \nu$ for a small $\nu > 0$, this can only cause exit when the receiver's type is in $[\theta_\infty, \theta_\infty + \nu]$ *and* the shock realization is in $[\eta - \nu, \eta]$, an event with probability bounded by a multiple of ν^2. The cost of taking this slight risk is thus second-order. The gain, on the other hand, is first-order: the proposer gains $\nu\pi'(x_t)$ conditional on still being in the game at that point.

Though this argument tells us that the proposer must at least take a *slight* risk that the receiver will exit in each period, it does not rule out major escalations. Indeed, moments of "crisis" characterized by a major increase in demands and high risk of exit may alternate with periods of slow escalation and a low, but positive, risk of immediate exit. The main factor in determining the shape of the proposer's optimal demand path is the shape of the distribution of receiver types, F.

Proposition 3. *Denote $\pi(x)(1 - F(x)) = u(x)$. Assume that π and F are such that (1) has finitely many solutions, i.e., u has finitely many critical points; $u''(x) \neq 0$ at all of them; and $u(x)$ takes different values at all of them. Say a critical point x^* of u is a stopping point if $u(x^*) > u(x')$ for all $x' > x^*$. Let $\theta_1^* < \ldots < \theta_k^*$ be the ordered set of stopping points. For each $1 \leq i \leq k-1$, let θ_i' be the (unique) value of $\theta \in \left(\theta_i^*, \theta_{i+1}^*\right)$ such that $u(\theta_i') = u(\theta_{i+1}^*)$.*

Take any sequence (η_n) with $\eta_n \searrow 0$, and G_n satisfying A1(η_n) for each n. Take any sequence of optimal demand paths $x^n = (x_t^n)_t$, where x^n is optimal when the shock distribution is G_n. Fix $\nu > 0$. Then, as $n \to \infty$, the proposer spends arbitrarily many periods making demands in any subinterval of $(\theta_i^* + \nu, \theta_i' - \nu)$ with positive measure, but never makes demands in $(\theta_i' + \nu, \theta_{i+1}^* - \nu)$, for all i.

The intuition behind Proposition 3 is as follows. Consider a proposer who was to play the unperturbed game, but against a modified distribution of receiver types—in particular, against a distribution of receiver types that was truncated so that only types above some threshold θ_0 remain. Results from Tirole (2016) tell us that the

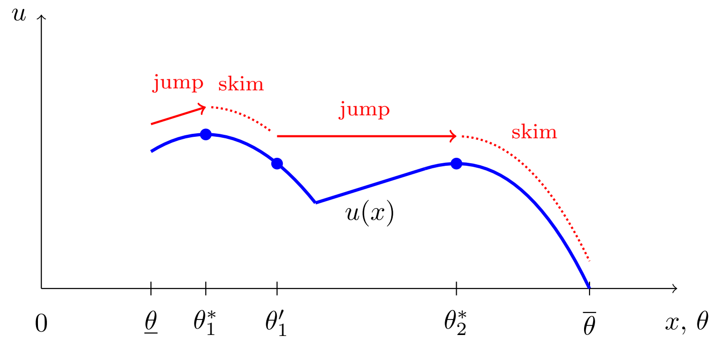

Figure 1: Pattern of escalation for small η and f s.t. $f(x) \equiv 0.5$ for $x \in [1, 2.5]$, $f(x) \equiv 0$ for $x \in [2.5, 3.5]$, $f(x) = (x - 3.5)/3$ for $x \in [3.5, 5]$

proposer's optimal demand path from there on is to jump to the demand x that maximizes u subject to the constraint $x \geq \theta_0$. That is, if $u' > 0$ at the bottom of the distribution, then the proposer wants to move her demand up to the lowest stopping point left in the distribution, which in general would constitute an instant, discontinuous escalation. If $u' < 0$, it is possible that the global maximum of u subject to the constraint $x \geq \theta_0$ is at θ_0, in which case the proposer would not screen out any receiver types at all.

Yet, when the receiver's behavior is noisy, the proposer will *always* be tempted to screen out at least a few marginal receivers in each period. Thus, when the proposer would not screen at all in the unperturbed game, she "nibbles" at the bottom of the distribution in the perturbed game. This is illustrated in the example given in Figure 1. In that example, the the proposer immediately screens out the types in the interval $[\underline{\theta}, \theta_1^*]$, then nibbles at the bottom of the distribution, which leads him to move her demand further and further away from the one-shot optimal demand θ_1^*. Eventually, enough of the distribution is nibbled away that the marginal receiver type crosses θ_1'; at this point, the proposer has screened out enough marginal receiver types that she might as well screen out a large set of additional receivers in one fell swoop, in order to move up to the next stopping point, θ_2^*. Afterwards, the proposer returns to slowly nibbling at the bottom of the remaining distribution.

In particular, suppose that F satisfies the (strict) monotone likelihood-ratio property (MHRP), that is, $\frac{1-F(x)}{f(x)}$ is a (strictly) decreasing function of x. Then $k = 1$, and the pattern of escalation is given by an initial demand close to the optimal one-shot demand θ_1^*, followed by gradual escalation for all $t > 0$, which becomes arbitrarily slow for η close to zero. If F does not satisfy this property, stopping points tend to be located in areas of high density; the proposer tends to nibble through such high-density parts of the distribution, and then escalate quickly through the valleys.

4 Full Commitment With Forward-Looking Receiver

Next, we show that, when the proposer has commitment power, the intuition from the case of a myopic receiver carries through to the general case, *so long as the receiver is strictly more impatient than the proposer*, even if only a little. More precisely:

Proposition 4. *Assume G satisfies A1(η) for some $\eta > 0$; $\delta_1 > \delta_2$; and the proposer has commitment power. Then:*

(i) If $T = \infty$, then, under any proposer-optimal demand path, the receiver eventually quits with probability 1.

(ii) For each $T < \infty$, let $(x_t^T)_{t\geq 0}$ be any proposer-optimal demand path in the game with horizon T. Then, along this sequence of demand paths, the probability that the receiver quits by the end of the game converges to 1 as $T \to \infty$.

Proposition 5. *If the proposer has commitment power, G satisfies A1(η) for some $\eta > 0$, and $\delta_1 \leq \delta_2$, then the receiver's exit probability is uniformly bounded away from 1, for any T.*

The intuition behind the result is as follows. When the receiver is not myopic, a proposer who considers marginally increasing a demand x_t stands to collect one gain and two types of costs. First, as in the myopic case, the proposer gains whenever the receiver was strictly willing to accept the (effective) demand $x_t + \epsilon_t$, and hence accepts even when this demand is slightly increased; and, conversely, the proposer loses whenever the receiver was marginal under the original demand, and is induced to exit by a small increase in x_t. Both of these effects accrue in period t. The second cost now faced by the proposer is that a receiver who is marginal in any period $s < t$ may alter her behavior *in period s*, if she expects to still be in the game by the end of period t with positive probability. Indeed, an increase in x_t lowers the receiver's continuation value in all periods before t, unless she expected to quit for sure by then. Hence, if she is "on the fence" in period s, there will be more realizations of the shock ϵ_s for which she chooses to quit.

The same logic as in Proposition 2 explains why the interplay of the first two forces eventually leads the proposer to force all receiver types to exit. Indeed, suppose that all types below a threshold θ_∞ eventually quit with probability 1, but those above θ_∞ remained forever with positive probability. For t large enough that (almost) all types below θ_∞ have quit, the gain from marginally increasing the demand x_t (so that types slightly above θ_∞ might now quit as well) exceeds the cost stemming from the risk that a marginal receiver might now quit in period t—or, for that matter, shortly before period t—simply because, by then, very few receiver types are left that are marginal given this demand. However, for large t, the second cost—that the receiver will "retaliate" by quitting preemptively in periods long before t—can potentially accrue over many periods. The assumption $\delta_2 < \delta_1$ ensures that this force washes out for large t: in most periods $s << t$, the receiver puts much lower weight on an expected loss in period t than the proposer puts on her potential gain. As a result, the change in the receiver's behavior in period s is also very small and negligible for the proposer. In contrast, when $\delta_2 \geq \delta_1$, the proposer stops short of escalating to the point where all receiver types quit, precisely because "proactive" retaliations by the receiver become too costly.

5 Extension: General Contracts (very preliminary!)

In the baseline model, we assume that the proposer can only commit to a sequence of demands $(x_t)_{t\geq 0}$. In principle, the proposer could be allowed to propose more general contracts, which give the receiver more freedom to vary her current transfer x_t without quitting the game, but at the cost of affecting the demands that the proposer will make in the future. We will see that, in this case, the receiver still quits the game eventually, although under more stringent conditions on the shock distribution.

A general contract takes the form $(X_t)_{t\geq 0}$, where $X_t(x_0, \ldots, x_{t-1}) \subseteq \mathbb{R}$ is defined for any feasible sequence $x_0, \ldots, x_{t-1}$. In words, a contract specifies a set X_0 of possible initial transfers x_0; a set $X_1(x_0)$ of allowed transfers in period 1 that in general depends on the initial transfer; a set $X_2(x_0, x_1)$ of allowed period 2 transfers, again conditioning on the history, and so on.

Intuitively, the proposer may want to make future demands a decreasing function of the receiver's current transfer: a high payment up front may be worthwhile for high types who expect to stay for a long time, but not for marginal types who will only choose to stay in the game today if they face a favorable shock.

We can prove the following result:

Proposition 6. *Assume G satisfies A2(η) for some $\eta > 0$; $T = \infty$; and $\delta_1 > \delta_2$. Suppose the proposer has commitment power and access to general contracts. Then the receiver eventually quits w.p. 1.*

This proposition is analogous to Proposition 4, but note that the result only holds when G satisfies A2(η), rather than A1(η).

6 No Commitment

We now consider the case of a proposer without commitment power, who faces a receiver that is not myopic ($\delta_2 > 0$). A naive intuition is that, in this case, the proposer will eventually induce all receiver types to exit, even under weaker conditions than when he has commitment power. In particular, unlike in Proposition 4, the condition $\delta_2 < \delta_1$ ought to be unnecessary: even if an escalation for large t will be retroactively punished by a higher likelihood of exit *before* period t, and even if the cost makes such an escalation unprofitable, ex post—once period t has been reached—the proposer will be tempted to escalate anyway, since the cost has already accrued.

This intuition, however, is incomplete: though the receiver's response before t to the *expected* value of x_t no longer affects the proposer's incentives at time t, the proposer now has to worry that a change in x_t may affect the receiver's expectations at time t about demands coming in periods *after* t. Indeed, it is possible that, if the proposer makes a slightly higher demand in period t that induces some additional receiver types to exit, his incentives in period $t+1$ will change in a way that induces a drastically higher equilibrium value for x_{t+1}. Expecting this, the receiver may respond harshly even to a tiny increase in x_t, which signals a much lower continuation value. This issue does not arise when the proposer has commitment power, since there is no danger of an unexpected increase in x_{t+1}—in particular, the proposer can commit ex ante both to a slightly higher value of x_t and to no change in x_{t+1}.

Here is a concrete example. Let $T = 1$, so that there are two periods. Assume f is such that, if no receiver types quit in period 0 (so $f_1(\theta; x) \equiv f_0(\theta) \equiv f(\theta)$) then, in period 1, the proposer is indifferent between two optimal demands x_*, x^*, with $x_* \in (\underline{\theta} - \eta, \underline{\theta} + \eta)$ and $x^* > \underline{\theta} + \eta$. (We can construct f with this property in the following way: set $f(\theta) = M_1$ for θ in a small right-neighborhood of M_1, $f(\theta) = M_2$ in a small neighborhood of some $\theta' \in (\underline{\theta}, \overline{\theta})$, and $f(\theta) = m$ elsewhere, with M_1, M_2 high and m low. We can then choose M_1 and M_2 so that the proposer is approximately indifferent between making a low demand that screens out almost nobody, and making a demand close to θ' that screens out all of the types near $\underline{\theta}$.)

It is easy to show that, if instead the proposer's initial demand x_0 induced *any* positive amount of screening of low types (so $f_1(\theta; x) < f_0(\theta; x)$ for some $\theta < \underline{\theta} + \eta$), then the proposer would strictly prefer to demand x^* in period 1. Thus the continuation utility expected by a low-type receiver in period 0 drops discontinuously when the initial demand x_0 crosses the threshold value $\underline{\theta} - \eta$, and hence so do the acceptance probabilities for low type receivers. The proposer may then prefer not to "nibble" at the bottom of the distribution at all in period 0, since an arbitrarily small transgression would lead to a nonzero mass of receiver types quitting.

Conjecture 1. Suppose that G satisfies A1(η), $T = \infty$, and the proposer has no commitment power. Then, if F satisfies the strict MHRP, then there is a PBE in which the receiver quits with probability 1.

Proposition 7. *Assume G satisfies A1(η) and $T = \infty$. Then, for any δ_1, $\delta_2 < 1$, the receiver quits with probability 1 in any smooth equilibrium.*

Hidden deviations

As we have seen, if the receiver can observe both the demand x_t and the idiosyncratic shock ϵ_t directly, then the proposer may refrain from inducing marginal receivers to exit, since a demand that risks exit may cause the receiver to expect much harsher demands in the future, and hence may induce more quitting today.

Since the receiver's payoff only depends on the sum $x_t + \epsilon_t$, however, we could alternatively assume that, in each period t, the receiver only observes the "effective demand" $y_t = x_t + \epsilon_t$. For a fixed demand path $(x_t)_t$, the receiver's best response is the same regardless of whether she observes x_t and ϵ_t separately, or only sees their sum. But assuming that only y_t is observed changes the receiver's response to deviations: if the proposer deviates by slightly increasing x_t relative to its equilibrium value, for instance to $x_t + \nu$ ($\nu > 0$), then the receiver in most cases will not realize that a deviation has occurred. Instead, she will chalk up the value of y_t to a slightly above-average shock realization.

Thus, the proposer can "hide" a deviation most of the time—more precisely, whenever $\epsilon_t \in [-\eta, \eta - \nu]$. For $\epsilon_t \in [\eta - \nu, \eta]$, equilibrium beliefs are not pinned down by Bayes' rule: both an off-path value of x_t and a shock realization $\epsilon_t \notin [-\eta, \eta]$ are zero-probability events. We will say that the receiver is *naive* if, upon seeing $y_t \notin [x_t - \eta, x_t + \eta]$, she still believes that the proposer offered x_t with probability 1, and that the shock realization is $\epsilon_t = y_t - x_t$.[4] Note that this assumption does *not* mean that the receiver no longer responds to a higher (effective) demand by rejecting more often; she still takes into account the increased cost of acquiescence. She does not, however, expect further deviations by the proposer; she believes the proposer is unaware that today's demand was high and will continue making her equilibrium demands in the future.

The following Proposition provides two conditions under which, when the receiver only observes effective demands, the proposer is again always tempted to risk exit by marginal receiver types, eventually leading to the end of the relationship.

Proposition 8. *Assume that $T = \infty$ and that the receiver observes only effective demands in each period. Then, if either*

(i) G satisfies A1(η) and the receiver is naive, or

(ii) G satisfies A2(η),

then, for any $\delta_1, \delta_2 \in [0, 1)$, in any PBE, the receiver eventually quits with probability 1.

The intuition is as follows. When the receiver is naive, the problem faced by the proposer is similar to the problem she faces under commitment—in particular, the receiver does not expect higher demands tomorrow upon seeing a higher demand today—but she no longer takes into account the "retroactive" impact of making a high demand in period t on the receiver's behavior in earlier periods. Thus, the argument made in Proposition 4 goes through regardless of the relationship between δ_1 and δ_2. Of course, for $\delta_2 \geq \delta_1$, the proposer's gradual escalation of demands is not in his best interest *ex ante*, but he is tempted to engage in this behavior *ex post*. Even for $\delta_2 < \delta_1$, the proposer's FOC pinning down the equilibrium values of

[4] This posterior belief would be uniquely selected if we considered an alternative model in which G had full support, and then we progressively reduced the weight G puts on shock realizations outside of $[-\eta, \eta]$ to zero in the limit.

$(x_t)_t$ differs from that obtained under commitment power, so the optimal demand path under full commitment is not time-consistent and the equilibrium demand path without commitment is *ex ante* inefficient.

Part (ii) shows that, if A2(η) is satisfied, then it does not matter for the result how the receiver interprets off-path realizations of y_t. The logic is as follows. Suppose, as before, that in equilibrium not all receiver types quit with probability 1, and more precisely, only types below a threshold θ_∞ do so. Consider the proposer's incentive to deviate by marginally increasing x_t by a small amount $\nu > 0$, for t large enough that the distribution of remaining receiver types is close to its limit. As usual, the direct benefit of this deviation is approximately $\pi'(x_t)F_{t+1}(\overline{\theta})\nu$, a positive multiple of ν. The cost is bounded approximately by an expression of the form $(1-G(\eta-\nu))M$, where M is a bound for the proposer's utility, and hence her net loss from causing a rejection (e.g., $M = \frac{\overline{\theta}+\eta}{1-\delta}$). The reason is that, with probabilty $G(\eta - \nu)$, the shock realization is low enough that the effective demand observed by the receiver is on-path—which, for t large enough, implies a probability of rejection arbitrarily close to zero. With the complementary probability, the receiver knows that an off-path event has occurred. Even if this inference leads to *all* receiver types quitting immediately, the proposer's cost is not more than M. Assumption A2(η) then implies that the benefit of a deviation dominates the cost for $\nu > 0$ small enough, as $g(\eta) = 0$ and hence $1 - G(\eta - \nu) \in o(\nu)$.

Ex-post observation of the shock

An alternative specification of the model allows the proposer to observe the realization of the shock ϵ_t, *but only at the end of period t*. Thus, the shock still serves the purpose of making the receiver's behavior slightly opaque from the proposer's point of view at the moment when he chooses his demand x_t. However, the proposer's posterior belief about the receiver's type is different from his belief in the baseline model: now, having observed ϵ_t, the proposer knows at the *end* of period t exactly which receiver types *would* have accepted a demand x_t. The posterior distribution $f_t(\theta)$ of receiver types at the beginning of period t thus takes the form $f(\theta)\mathbb{1}_{\theta \geq \theta_t^*}$ for some threshold type θ_t^*.

An interpretation of this setting is as follows. Suppose that x_t is a task that the principal requests an employee to complete, yielding utility $\pi(x_t)$ for the principal.

The effective difficulty of completing the task, $x_t+\epsilon_t$, depends on idiosyncratic factors that are unpredictable from the principal's point of view. However, if the employee does accept and she completes the task, the principal may *ex post* learn about the hurdles involved and hence how much effort the employee had to put in to complete the task.

Since both players have common knowledge of the proposer's belief about the distribution of remaining receivers at the beginning of any period t, we can study Markov perfect equilibria (MPE). At the beginning of each period t, the relevant state variable is the lowest type of the receiver θ_t^* that would still be in the game, given equilibrium strategies. Moreover, when the receiver sees the proposer's demand and shock realization (x_t, ϵ_t respectively), it is only the effective demand $y_t = x_t + \epsilon_t$ and the incumbent threshold type θ_t^* that are payoff-relevant, the latter being relevant only if it is high enough that no receiver types quit today. Thus we can focus on equilibria in which the receiver uses a threshold strategy $\theta^*(y)$, accepting an effective demand y_t iff $\theta \geq \theta^*(y_t)$; the proposer's equilibrium demand x_t in period t takes the form $x^*(\theta_t^*)$; and the threshold type in the next period after an effective demand y is $\theta_{t+1}^* = \max(\theta_t^*, \theta^*(y))$. Moreover, we will focus on equilibria in which the functions θ^* and x^* are weakly increasing. For simplicity we refer to MPE with these additional properties simply as equilibria.

We then have the following

Proposition 9. *Assume that G satisfies A2(η) and $T = \infty$. For any $\delta_1, \delta_2 \in [0, 1)$, in any equilibrium, the receiver eventually quits with probability 1.*

Moreover, if G only satisfies A1(η), the same result holds in any continuous *equilibrium, that is, any equilibrium in which the mappings x^* and θ^* are continuous.*

The intuition is as follows. Much as in Proposition 8.(ii), with probability $G(\eta-\nu)$, a deviation of size ν at a large t would have no effect, since the shock realization would be small enough that $\theta^*(y_t) \leq \theta_t^*$. The proposer can thus "get away with" such a small deviation with probability $1 - o(\nu)$. The difference with Proposition 8 is that here we do not need to assume that the deviation is hidden—it makes no difference that the receiver directly observes the proposer's intended demand x_t, since all equilibrium behavior from the receiver's decision in period t onwards conditions only on the effective demand y_t.

The full-commitment setting (and the special case of a fully myopic receiver) can also be studied in this alternative model. The results are qualitatively similar to those obtained in the baseline model. More precisely, with a myopic receiver, the commitment solution is time-consistent, and leads to sure exit under assumption A1(η); with a forward-looking receiver, time consistency breaks down, and the commitment solution leads to sure exit under assumption A1(η) if $\delta_2 < \delta_1$.

It is worth making one more observation on Propositions 8 and 9. It can be shown in the setting of Proposition 8 that, if G satisfies A1(η), then the receiver eventually quits with probability 1, even if the receiver directly observes the proposer's demands, in any equilibrium that is smooth in the following sense: letting $h^t = (x_0, \ldots, x_t)$ be a history of demands up to time t, $x_{t+1}(h^t) = x_{t+1}(x_0, \ldots, x_t)$ is a differentiable function of x_0, ..., x_t for all t. In other words, if when the proposer increases a demand x_t slightly, it only makes the receiver expect that other future demands might increase slightly as well, then the proposer is always willing to provoke some marginal types just as when the receiver is myopic; it is only when x_{t+1} might be expected to increase much more than x_t that the receiver might react disproportionately to an escalation, thus dissuading the proposer from escalating. Whether equilibria of this game are generally smooth is an open question. Similarly, in the setting of Proposition 9, even if only A1(η) is satisfied, it still holds that the receiver eventually quits with probability 1, in any equilibrium that is continuous in the sense that the function $x^*(\theta)$ is continuous in θ (which implies that $\theta^*(y)$ is also a continuous function). Again, whether equilibria in that setting are continuous in general is an open question.

Transition dynamics with ex post observable shocks (very preliminary!)

Consider for simplicity the case $\delta_2 = 0$. When the proposer observes the shock ϵ_t at the end of each period t, we can get a more explicit characterization of the speed at which the proposer "nibbles" away the bottom of the distribution of receivers.

In particular, we can show that, if $x^*(\theta) \in (\theta - \eta, \theta + \eta)$ and η is "small", then

$$G(\theta - x^*(\theta)) \approx \frac{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)} - 1}{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)} - \delta_1}. \tag{3}$$

Note that $G(\theta - x^*(\theta)) = P(\epsilon < \theta - x^*(\theta))$ is the equilibrium probability that the state does not change, i.e., the probability that the effective offer is such that no remaining receiver types *would* have quit. $1-G(\theta-x^*(\theta))=\int_{\theta-x^*(\theta)}^{\eta} g(\epsilon)d\epsilon \approx \frac{1-\delta_1}{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1}$ is the probability that the state moves upward by any amount. The probability that the receiver quits in the current period (conditional on still being in the game and on the current state) is

$$\int_{\theta-x^*(\theta)}^{\eta} F_\theta(x+\epsilon)g(\epsilon)d\epsilon \approx \int_{\theta-x}^{\eta} \frac{f(\theta)}{1-F(\theta)}(x+\epsilon-\theta)g(\epsilon)d\epsilon =$$
$$= \frac{f(\theta)}{1-F(\theta)}\int_{\theta-x^*(\theta)}^{\eta}(x^*(\theta)+\epsilon-\theta)g(\epsilon)d\epsilon = \frac{f(\theta)}{1-F(\theta)}\int_{\theta-x^*(\theta)}^{\eta}(1-G(\epsilon))d\epsilon$$

Now suppose $g(\epsilon)\approx(\eta-\epsilon)^k\tilde{g}$ for ϵ close to η, and $\theta-x$ is close to η. Then $1-G(\theta-x)\approx\frac{(\eta-\theta+x)^{k+1}}{k+1}\tilde{g}$, and the probability that the receiver quits is approximately $\frac{f(\theta)}{1-F(\theta)}\frac{(\eta-\theta+x)^{k+2}}{(k+1)(k+2)}\tilde{g} = \frac{f(\theta)}{1-F(\theta)}\frac{\eta-\theta+x}{k+2}(1-G(\theta-x))$. In particular, if $g\equiv\tilde{g}=\frac{1}{2\eta}$ is uniform, then $1-G(\theta-x)\approx(\eta-\theta+x)\tilde{g}$, and the probability that the receiver quits is $\approx\frac{f(\theta)}{1-F(\theta)}\frac{(\eta-\theta+x)^2}{2}\tilde{g}$, which we can calculate as follows:

$$1-G(\theta-x)\approx(\eta-\theta+x)\tilde{g}\approx\frac{1-\delta_1}{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1}$$
$$\eta-\theta+x\approx\frac{1}{\tilde{g}}\frac{1-\delta_1}{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1}$$
$$\frac{f(\theta)}{1-F(\theta)}\frac{(\eta-\theta+x)^2}{2}\tilde{g}\approx\frac{1}{2\tilde{g}}\frac{f(\theta)}{1-F(\theta)}\frac{(1-\delta_1)^2}{\left(\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1\right)^2}$$
$$\approx\eta\frac{f(\theta)}{1-F(\theta)}\frac{(1-\delta_1)^2}{\left(\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1\right)^2}$$

We can draw several takeaways. As δ_1 goes to 1, the proposer's offers become more generous. If we imagine this increase in the discount factor representing a world with more frequent offers, then, as $\delta_1\to 1$, the probability that the state moves up in any unit of real time remains roughly constant, since the probability that the state moves up due to any single offer, $1-G(\theta-x)$, is roughly proportional to $1-\delta_1$. Moreover, since the probability of exit is a multiple of $\frac{\eta-\theta+x}{k+2}(1-G(\theta-x))$, and $x\to\theta-\eta$ as $\delta_1\to 1$, the probability of exit per unit of real time actually decreases as offers become more frequent, though eventual exit remains inevitable for all $\delta_1<1$.

As $\eta\to 0$, $1-G(\theta-x)$ remains roughly constant, but the expected size of a jump shrinks (assuming that we compress G so that the shock is multiplied by a factor less than 1, the distribution of the jump size $x + \epsilon - \theta$ is compressed in the same way).

Transition dynamics with ex post observable shocks and patient receiver (very preliminary!)

We can obtain a solution in closed form, even for $\delta_2 > 0$, in the case where shocks are ex post observable, for the following parametric example. Assume that $F(\theta) = 0$ for $\theta < \underline{\theta}$ and $F(\theta) = 1 - e^{a(\theta - \underline{\theta})}$ for $\theta \geq \underline{\theta}$. In addition, assume that shocks are multiplicative, i.e., if the proposer makes a demand x_t in period t, the receiver's flow payoff from accepting is $-x_t - x_t \epsilon_t$, where $\epsilon_t \sim G$. Finally, assume that $\pi(x) \equiv x$. These assumptions are outside of the baseline model, where the support of F is bounded and shocks are additive and independent of the proposer's demands, but we can extend the model to accommodate this case and in fact the main results would extend.

Under these assumptions, there is an equilibrium of the game with the following structure: the proposer's strategy takes the form $x^*(\theta) \equiv x_0 \theta$ for all θ, and the receiver's strategy takes the form $\theta^*(y) \equiv y(1 - \omega)$ for all y, where x_0 and ω are fixed parameters. Moreover, $x_0 > \frac{1}{1+\eta(1-\delta_2)}$, so that it is optimal for the marginal receiver to quit with positive probability, and $\omega \geq 0$, since type θ will always accept if $y \leq \theta$.

A characterization of x_0 and ω are given in the Appendix. Let $z_0 = x_0(1 - \omega)$, and let ϵ^* be the marginal value of ϵ for which no receivers quit in that period, i.e., $x_0 \theta (1 + \epsilon^*)(1 - \omega) = \theta \Longrightarrow \epsilon^* = \frac{1}{z_0} - 1$.

As shown in the Appendix, we can derive an equation that pins down z_0, and in which all expressions are independent of δ_2. It follows that, as δ_2 varies, $x_0(1-\omega)$ stays constant; $\omega \equiv \omega_2 \delta_2$ for some ω_0 which is not a function of δ_2; and $x_0 = \frac{z_0}{1-\omega} = \frac{z_0}{1-\delta_2\omega_0}$. Intuitively, as δ_2 increases, the receiver becomes somewhat *more lenient* because she puts higher weight on the option value from staying in the game. This, paradoxically, allows the proposer to make *more aggressive* demands, but in such a way that ϵ^* (hence the rejection probability, and the evolution of the state) remain unchanged.[5] This is in stark contrast to our results in the commitment case, where a more patient receiver dissuades escalations by the proposer.

[5] In particular, the evolution of the state and the receiver's probability of exit in any period are exactly equal to those from the special case in which the receiver is completely impatient, i.e., $\delta_2 = 0$.

7 Extensions

Richer Receiver Actions (very preliminary!)

Our baseline model is quite stark: the receiver has to either end the game, or else let the proposer do as he will. Anecdotally, "salami tactics" are sometimes countered by a show of force that falls short of ending the game. For example, in the context of international negotiations, the receiver might respond to an escalation by expelling the proposer's ambassadors, imposing sanctions, etc.

What happens if we allow intermediate responses by the receiver? Intuitively, two effects arise. First, the receiver can now "act tough", so signaling concerns may induce more aggressive receiver behavior. (In the baseline model, the only way to "act tough" is to quit, so a receiver can never benefit in the continuation by exaggerating how low her type is.) Second, in equilibrium, more information about the receiver's type may be revealed. As a result, the proposer may be able to better tailor her demands to the receiver's true type, leading to less exit. We will see that, if the set of receiver actions is coarse (in particular, finite), then this latter effect cannot be strong enough as to overturn our main results.

We operationalize "intermediate" actions in a simple way. We now assume that the receiver has access to a finite set of quitting probabilities $0 = p^0 < \ldots < p^k = 1$ ($k \geq 2$). In each period t, after seeing the demand x_t and the shock ϵ_t, the receiver can choose any p^i, rather than staying for sure ($p = 0$) or exiting ($p = 1$) as in the baseline model. If p^i is chosen, then with probability p^i the game ends and the receiver realizes her terminal payoff; with probability $1 - p^i$, the players receive the flow payoffs $(\pi(x_t), -x_t - \epsilon_t)$ and the game continues. Crucially, we assume that the proposer sees the chosen p^i, not just its realization. Thus, a receiver who "rolls the dice" can signal toughness even if the outcome ends up being peace. (In contrast, if only the outcome of the receiver's action were observable, then choosing $p \in (0, 1)$ would just amount to mixing in the baseline model.) For brevity, we take $T = \infty$ throughout.

Proposition 10. *Propositions 4, 8 and 9 extend to this setting. That is:*

(i) Assume commitment power; A1(η) for some $\eta > 0$; and $\delta_1 > \delta_2$. Then, under any optimal demand path, the receiver eventually exits w.p. 1.

(ii) *Assume no commitment power; A2(η) for some $\eta > 0$; and the receiver observes only effective demands in each period. Then, for any $\delta_1, \delta_2 \in [0, 1)$, in any PBE, the receiver eventually quits w.p. 1.*

(iii) *Assume no commitment power and A2(η) for some $\eta > 0$ in the "ex post observed shocks" setting. Then, for any $\delta_1, \delta_2 \in [0, 1)$, in any equilibrium, the receiver eventually quits w.p. 1.*

Note that, in this variant of the model, the receiver may eventually exit w.p. 1 either by quitting outright ($p = 1$) once or by rolling the dice ($p > 0$) infinitely many times.

8 Discussion

The family of models we study represents settings in which the receiver is subject to a ratchet effect conditional on accepting the proposer's demands and thus remaining in the game. We have shown that a small perturbation to the problem, which makes the receiver's decisions slightly opaque from the proposer's point of view, can dramatically change the predictions of the model: when the receiver's decisions are unaffected by noise or this noise is transparent *ex ante* to the proposer (Tirole, 2016), the proposer generally chooses not to ratchet at all, whether or not she has commitment. In contrast, even a small amount of noise tempts the proposer to take small risks of rejection, which then snowball: once the proposer is quite confident that a given type θ of the receiver would most likely have quit by now, she is tempted to then risk rejection by types $\theta + \nu$ for small ν, and so on.

As we have seen, when the receiver is forward-looking, this simple logic is complicated in very different ways depending on whether the proposer has commitment power or not. When he does, the proposer must account for both the contemporaneous risk of causing exit due to escalation and the retroactive risk that a late escalation will prompt preemptive exit by the receiver in earlier periods, an effect that may limit escalation but only if the receiver is at least as patient as the proposer. When he does not, the proposer must instead account for what the receiver can observe if the proposer deviates; what she can then infer about possible changes in future demands; and how her behavior today will change as a result. When the receiver can directly see the proposer's intended demands (and hence can observe any deviations perfectly), the cost of even a small deviation can be high if it is seen as the harbinger of much more severe escalations in the future. When the receiver only sees the noisy outcome of the proposer's choice, it is much less likely that the proposer can be punished harshly for a small deviation, and hence the temptation to gradually escalate creeps back in.

The underlying intuition resonates with our informal understanding of salami tactics and how a prospective perpetrator of such tactics might be deterred, dating back to Schelling (1966). The receiver would like to convince the proposer that she will uphold a *bright line* x_0, reacting harshly if it is crossed ($x_t > x_0$). (For example: using nuclear or chemical weapons; crossing the Rubicon, etc.) Since our receiver is assumed not to have commitment power, such a threat is sustainable only when, in equilibrium, a crossing of the line signals that the proposer will engage in much worse transgressions in the future. When the proposer has commitment power, the notion of a bright line has no bite, since the proposer can commit to all future actions, and thus has full control over what the receiver might expect after a current demand that is perceived as too high. When the proposer lacks commitment, bright lines may work, but only when the receiver can see exactly what the proposer intended to demand. When any noise is involved at this stage, the proposer is able to blur any bright line the receiver might like to uphold, as willful escalations become confused with accidents.

A next step for this project is to give a sharper description of the transition dynamics generated by the proposer's escalation, at least under some conditions. A broader aim is to think more generally about how persistent private information may be subject to leakage when players' actions are observed by others with even a small amount of noise.

A Proofs

Proof of Proposition 2. That the receiver exits with probability 1 follows from the more general result in Proposition 4. But note that $F_t(\overline{\theta})$ must be positive for all t: if not, and t_0 is the lowest t for which $F_t(\overline{\theta}) = 0$, the proposer can do strictly better with any offer that is accepted with positive probability in period $t-1$.

Suppose for the sake of contradiction that $\liminf_{t\to\infty} x_t = \overline{\theta} - \eta - \nu$, where $\nu > 0$, and let $(x_{t_l})_l$ be a subsequence such that $x_{t_l} \xrightarrow[l\to\infty]{} \overline{\theta} - \eta - \nu$. Because, for all l large enough, $P_{t_l,t_l}(\theta) = 1$ for all $\theta > \overline{\theta} - \nu$, whereas $P_{t_l,t_l}(\theta) \to G(\theta - \overline{\theta} + \eta + \nu) < 1$ for all $\theta < \overline{\theta} - \nu$ (and of course $P_{t,t}(\theta)$ is weakly increasing in θ for all other t), we must have that, for any $\theta < \overline{\theta} - \eta - \nu < \theta'$, either $P_t(\theta) = 0$ for t large enough, or $\frac{P_t(\theta)}{P_t(\theta')} \xrightarrow[t\to\infty]{} 0$. For l large enough, the optimality of x_{t_l} then contradicts (2). This finishes the proof of part (ii).

For part (i), it is easy to show that the proposer's problem is time consistent, since adjusting $(x_t)_{t\geq s}$ to optimize his continuation value at time s does not affect the receiver's behavior in periods $0, 1, \ldots, s-1$. $\square$

Proof of Proposition 3. only a sketch, needs cleanup

We write the proof for the case in which $u'(\underline{\theta}) > 0$, i.e., it is optimal to screen some receiver types in the unperturbed game; the alternative case is similar.

We first show the second part. Suppose that, for some i, there is a subsequence $n_k \to \infty$, such that for all k there is a demand $x^{n_k}_{t_{n_k}} \in (\theta'_i + \nu, \theta^*_{i+1} - \nu)$. For each k, take t_{n_k} to be minimal with this property. In addition, without loss of generality, and by an application of Bolzano-Weierstrass, assume that the subsequence $(n_k)_k$ is such that $x^{n_k}_{t_{n_k}}$ converges to a limit $x^* \in [\theta'_i + \nu, \theta^*_{i+1} - \nu]$, and $x^{n_k}_{t_{n_k}-1}$ converges to a limit $x^{**} \leq \theta'_i + \nu$.

Since the proposer's problem is time-consistent when $\delta_2 = 0$, if the demand paths x^{n_k} are optimal, in particular the continuation demand path $(x^{n_k}_t)_{t\geq t_{n_k}}$ must be optimal for each k given the residual distribution of receivers $F_{t_{n_k}}(\theta; x^{n_k})$. Denote by $U(x; H; n)$ the proposer's utility from a demand path x when the initial receiver type distribution is H and $G = G_n$. Then, for any k, $U((x^{n_k}_t)_{t\geq t_{n_k}}; F_{t_{n_k}}(\cdot; x^{n_k}); n_k)$ must be weakly higher than what the proposer can obtain from any other demand path, in particular, from the demand path $y = (y_t)_t$ given by $y_t \equiv y^*$, where

$y^* = \arg\max_{x \geq x^{**}} u(x)$. That is, for all k, we must have

$$U((x_t^{n_k})_{t \geq t_{n_k}}; F_{t_{n_k}}(\cdot; x^{n_k}); n_k) \geq U(y; F_{t_{n_k}}(\cdot; x^{n_k}); n_k). \tag{4}$$

However, as $k \to \infty$, the right-hand side of (4) converges to the maximal utility the proposer can obtain in the unperturbed game. Indeed, by construction, y^* is the optimal one-shot demand when facing a receiver distribution $\propto f(\theta)\mathbb{1}_{\theta \geq x^{**}}$, and the optimality of repeating the one-shot demand forever follows from Tirole (2016). Moreover, y^* must equal either θ^*_{i+1} or θ^*_j for some $j \leq i$ (in particular if $x^{**} \leq \theta^*_i$); either way, $y^* \notin (\theta'_i, \theta^*_{i+1})$. On the other hand, the limsup of the left-hand side of (4) is at most what the proposer could obtain in the unperturbed game from an optimal demand path *subject to the constraint that the first demand must be x^**. It can be shown that this payoff is strictly lower than that obtained from the demand path y [[[details]]], so (4) must fail to hold for large enough k, a contradiction.

Next, we show the first part. It is enough to show that, given any interval $(a, b) \subseteq (\theta^*_i + \nu, \theta'_i - \nu)$ with $a < b$, the proposer makes at least one demand in (a, b) for all n large enough. The result can then be obtained by considering collections of disjoint open intervals contained in $(\theta^*_i + \nu, \theta'_i - \nu)$ with arbitrarily many elements.

For the sake of contradiction, suppose that there is i, ν, $a < b$, and a subsequence $n_k \to \infty$ such that $x^{n_k} \cap (a, b) = \emptyset$ for all k. For each k, let t_{n_k} be the lowest t for which $x_t^{n_k} \geq b$. (For all $k \geq k_0$, such a t_{n_k} must exist by Proposition 2.) Then, by construction, $x^{n_k}_{t_{n_k}-1} \leq a$. Again, by Bolzano-Weierstrass, we can assume that $x^{n_k}_{t_{n_k}}$ converges to a limit $x^* \geq b$, and that $x^{n_k}_{t_{n_k}-1}$ converges to a limit $x^{**} \leq a$. Then, for all k, we must have

$$U((x_t^{n_k})_{t \geq t_{n_k}}; F_{t_{n_k}}(\cdot; x^{n_k}); n_k) \geq U(z; F_{t_{n_k}}(\cdot; x^{n_k}); n_k),$$

where $z = (z_t)_t$ is given by $z_t \equiv z^*$, where $z^* = \arg\max_{x \geq x^{**}} u(x)$. Here, by construction, z^* must be either θ^*_j for some $j \leq i$ (if $x^{**} \leq \theta^*_i$) or x^{**} (if not). In particular, in the unperturbed game where the proposer faces a receiver distribution $\propto f(\theta)\mathbb{1}_{\theta \geq x^{**}}$, the optimal demand path never demands more than $\max(\theta^*_i, x^{**})$, so a path with initial demand $x^* \geq b$ is strictly suboptimal, and again the above inequality fails for large enough k. □

Proof of Proposition 4. We will first prove part (i).

As in Proposition 2, define $U_t(\theta; x)$ as the proposer's continuation utility at time t, if the receiver is of type θ and is still in the game; $P_t(\theta; x)$ as the probability that a receiver of type θ accepts through period t, given an equilibrium demand path x; let $f_t(\theta; x) = f(\theta)P_{t-1}(\theta; x)$ be the density of receiver types faced by the proposer in period t; and $F_t(\theta; x) = \int_{-\infty}^{\theta} f_t(\tilde{\theta}; x)d\tilde{\theta}$. More generally, for $t \geq s$, let $P_{t,s}(\theta; x)$ be the probability that a receiver of type θ accepts through period t, conditional on being in the game at the beginning of period s. In particular, $P_t(\theta; x) = P_{t,0}(\theta; x)$, and $P_{t,t}(\theta; x)$ is the probability that this receiver accepts in period t.

Let $V_t(\theta; x)$ be the continuation payoff of a receiver of type θ at the beginning of period t, conditional on a demand path x. For general T,

$$\begin{aligned} V_t(\theta; x) &= E_{\epsilon_t}\left[V_t(\theta; x, \epsilon)\right], \text{ where} \\ V_t(\theta; x, \epsilon) &= \max\left[-x_t - \epsilon_t + \delta_2 V_{t+1}(\theta; x), -\frac{1-\delta_2^{T-t}}{1-\delta_2}\theta\right]. \end{aligned} \tag{5}$$

By the envelope theorem,

$$\begin{aligned} \frac{\partial V_t(\theta; x)}{\partial \theta} &= -\sum_{T \geq s \geq t} \delta_2^{s-t}(1 - P_{s,t}(\theta)) \\ \text{and, for } s \geq t, \quad \frac{\partial V_t(\theta; x)}{\partial x_s} &= -\delta_2^{s-t} P_{s,t}(\theta). \end{aligned} \tag{6}$$

By (5), a receiver of type θ accepts at time t iff $\frac{1-\delta_2^{T-t}}{1-\delta_2}\theta \geq x_t + \epsilon_t - \delta_2 V_{t+1}(\theta; x)$, i.e., she accepts with probability

$$P_{t,t}(\theta; x) = G\left(\frac{1-\delta_2^{T-t}}{1-\delta_2}\theta - x_t + \delta_2 V_{t+1}(\theta; x)\right). \tag{7}$$

The marginal change in her probability of acceptance at time t, if x_s is increased marginally for $s \geq t$, is then

$$\frac{\partial P_{t,t}(\theta; x)}{\partial x_s} = \begin{cases} -g\left(\frac{1-\delta_2^{T-t}}{1-\delta_2}\theta - x_t + \delta_2 V_{t+1}(\theta; x)\right) \text{ if } s = t \\ -g\left(\frac{1-\delta_2^{T-t}}{1-\delta_2}\theta - x_t + \delta_2 V_{t+1}(\theta; x)\right) \delta_2^{s-t} P_{s,t+1}(\theta) \text{ if } s > t, \end{cases} \tag{8}$$

where we have used that $\frac{\partial V_{t+1}(\theta; x)}{\partial x_s} = -\delta_2^{s-t-1} P_{s,t+1}(\theta)$.

The proposer's problem, as in Proposition 2, is

$$\max_x \sum_{t=0}^{T} \delta_1^t \pi(x_t) \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_t(\theta; x) d\theta.$$

For each $s \geq 0$, the solution must satisfy the FOC:

$$0 = \frac{\partial E(U_0(\theta; x))}{\partial x_s} = \delta_1^s \pi'(x_s) \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_s(\theta; x) d\theta +$$

$$\sum_{t=0}^{s} \delta_1^t \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{t-1}(\theta; x) \frac{\partial P_{t,t}(\theta; x)}{\partial x_s} \overline{U}_t(\theta; x) d\theta$$

$$= \delta_1^s \pi'(x_s) \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_s(\theta; x) d\theta -$$

$$\sum_{t=0}^{s} \delta_1^t \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{t-1}(\theta; x) g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right) \delta_2^{s-t} P_{s,t+1}(\theta) \overline{U}_t(\theta; x) d\theta$$

$$\pi'(x_s) F_{s+1}(\overline{\theta}; x) =$$

$$\sum_{t=0}^{s} \left(\frac{\delta_2}{\delta_1}\right)^{s-t} \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{t-1}(\theta; x) g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right) P_{s,t+1}(\theta) \overline{U}_t(\theta; x) d\theta, \tag{9}$$

where we have substituted in $T = \infty$, and $\overline{U}_t(\theta; x) = x_t + \delta_1 U_{t+1}(\theta; x)$ is the proposer's continuation value at time t conditional on the receiver accepting in period t. Note that the FOC for each s must hold with equality—in particular, we cannot have a corner solution—for the following reasons. A proposal $x_t > \frac{\overline{\theta}+\eta}{1-\delta_2}$ will always be rejected by the receiver, and is thus strictly suboptimal for the proposer. In addition, because $\underline{\theta} > \eta$, a small enough positive proposal $x_s \in [0, \underline{\theta} - \eta)$ can never be rejected by the receiver. Making a proposal $x_s < \underline{\theta} - \eta$—in particular, making a proposal $x_s = 0$—could then be optimal only because it might increase the probability that the receiver accepts in periods before s. But, if so, shifting x_s up by some small $\nu > 0$ and shifting x_{s-1} down by $\delta_2 \nu$ leaves the receiver's acceptance constraint at time $s-1$ (hence also at earlier times) unchanged, and is a strict improvement for the proposer. Then a proposal $x_s = 0$ can be part of an optimal demand path only if $x_t = 0$ for all $t < s$. In particular, we must have $x_0 = 0$, which can never be optimal. Thus the optimal demand in every period must lie in the open interval $\left(0, \frac{\overline{\theta}+\eta}{1-\delta_2}\right)$.

Let $P_t(\theta;x) \searrow P_\infty(\theta;x)$, $F_t(\theta;x) \searrow F_\infty(\theta;x)$. Assume the receiver does not quit with probability 1, so that $F_\infty(\overline{\theta};x) > 0$.

As we noted, proposals $x_t > \frac{\overline{\theta}+\eta}{1-\delta_2}$ are suboptimal, so $x_s \leq \frac{\overline{\theta}+\eta}{1-\delta_2} := \overline{x}$ for all s. Then $\pi'(x_s) \geq \pi'(\overline{x}) > 0$ for all s. Note that $\overline{U}_t(\theta;x) \leq \frac{\overline{\theta}+\eta}{(1-\delta_1)(1-\delta_2)} := M$, and $P_{\cdot,\cdot}(\theta;x) \leq 1$. Let $\overline{f} = \sup_\theta f(\theta) < \infty$ and $\overline{g} = \sup_\epsilon g(\epsilon) < \infty$.

Take t_0 such that $\left(\frac{\delta_2}{\delta_1}\right)^{t_0} \frac{\delta_1}{\delta_1-\delta_2}\overline{g}M < \frac{1}{2}\pi'(\overline{x})F_\infty(\overline{\theta};x)$. In particular, $\sum_{t=t_0}^{s}\left(\frac{\delta_2}{\delta_1}\right)^t \overline{g}M < \frac{1}{2}\pi'(\overline{x})F_\infty(\overline{\theta};x)$ for all $s \geq t_0$.

Note now that (9) implies that, for all s,

$$\begin{aligned}
\pi'(\overline{x})F_\infty(\overline{\theta};x) \leq &\sum_{t=0}^{s-t_0}\left(\frac{\delta_2}{\delta_1}\right)^{s-t}\int_{\underline{\theta}}^{\overline{\theta}} f(\theta)\overline{g}M d\theta \\
&+ \sum_{t=s-t_0+1}^{s}\left(\frac{\delta_2}{\delta_1}\right)^{s-t}\int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)P_{s,t+1}(\theta)\overline{U}_t(\theta;x)d\theta \\
\frac{1}{2}\pi'(\overline{x})F_\infty(\overline{\theta};x) < &\sum_{t=s-t_0+1}^{s}\int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)P_{s,t+1}(\theta)\overline{U}_t(\theta;x)d\theta \\
\frac{1}{2}\pi'(\overline{x})F_\infty(\overline{\theta};x) < &\sum_{t=s-t_0+1}^{s}\int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)M d\theta
\end{aligned}$$

In particular, there are arbitrarily high values of t for which

$$\frac{1}{2Mt_0}\pi'(\overline{x})F_\infty(\overline{\theta};x) < \int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)d\theta.$$

Let $\theta_\infty = \inf(\mathrm{supp}(P_\infty(\theta;x)))$. For $\theta < \theta_\infty$, $f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right) \leq \overline{f}\overline{g}$, and this expression converges pointwise to zero as $t \to \infty$, since $P_{t-1}(\theta) \to 0$ and f, g are bounded. Then, by the dominated convergence theorem,

$$\int_{\underline{\theta}}^{\theta_\infty} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)d\theta \xrightarrow[t\to\infty]{} 0.$$

Then there must be arbitrarily high values of t for which

$$\frac{1}{2Mt_0}\pi'(\overline{x})F_\infty(\overline{\theta};x) \leq \int_{\theta_\infty}^{\overline{\theta}} f(\theta)P_{t-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta;x)\right)d\theta. \quad (10)$$

Recall that type θ quits at time t iff $\frac{\theta}{1-\delta_2} < x_t + \epsilon_t - \delta_2 V_{t+1}(\theta; x)$. Thus, if $\frac{\theta}{1-\delta_2} < x_t + \eta - \nu - \delta_2 V_{t+1}(\theta; x)$ for some t and $\nu > 0$, type θ would quit with probability $1 - G(\eta - \nu) > 0$ at that time. Hence, if there is a $\nu > 0$ for which there are arbitrarily high values of t such that $\frac{\theta}{1-\delta_2} < x_t + \eta - \nu - \delta_2 V_{t+1}(\theta; x)$, then type θ would quit with probability 1. Thus, for any $\theta > \theta_\infty$, $\liminf_{t\to\infty} \frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x) \geq \eta$. By (6), $V_{t+1}(\theta; x)$ is continuous in θ, so this implies $\liminf_{t\to\infty} \frac{\theta_\infty}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta_\infty; x) \geq \eta$.

Note also that $\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)$ is a weakly increasing function of θ, by (6). Thus, if $\theta \in \operatorname{supp}(P_\infty(\theta; x))$, and $\theta' \geq \theta$, then $\theta' \in \operatorname{supp}(P_\infty(\theta; x))$, implying that $(\theta_\infty, +\overline{\theta}) \in \operatorname{supp}(P_\infty(\theta; x))$. Moreover, by (6), and the fact that $P_{s,t}(\theta)$ is weakly increasing in θ, $\frac{\partial V_t(\theta;x)}{\partial \theta}$ is weakly increasing in θ. In addition, for any type $\theta > \theta_\infty$, we must have $\frac{\partial V_t(\theta;x)}{\partial \theta} \xrightarrow[t\to\infty]{} 0$, again by (6). It follows that, for any $\nu > 0$, the derivative of $\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)$ with respect to θ is positive and uniformly bounded away from zero for all $(t, \theta) \in \{t_1, t_1 + 1, \ldots\} \times (\theta_\infty + \nu, \overline{\theta})$, for some t_1 large enough.

By (10), there are arbitrarily high values of t for which

$$
\begin{aligned}
&\frac{1}{2Mt_0}\pi'(\overline{x})F_\infty(\overline{\theta}; x) \leq \overline{f} \int_{\theta_\infty}^{\overline{\theta}} g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right) d\theta \\
\implies &\frac{1}{2Mt_0}\pi'(\overline{x})F_\infty(\overline{\theta}; x) \leq \overline{f}\overline{g} |\{\theta \geq \theta_\infty : \frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x) \leq \eta\}|.
\end{aligned}
$$

Take $\nu_0 \in \left(0, \frac{\pi'(\overline{x})F_\infty(\overline{\theta};x)}{4Mt_0\overline{f}\overline{g}}\right)$. Then we must have that, for arbitrarily large t,

$$
\frac{1}{4Mt_0\overline{f}\overline{g}}\pi'(\overline{x})F_\infty(\overline{\theta}; x) \leq |\{\theta \geq \theta_\infty + \nu_0 : \frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x) \leq \eta\}|.
$$

But, since $\liminf_{t\to\infty} \frac{\theta_\infty}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta_\infty; x) \geq \eta$, and the derivative of $\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)$ with respect to θ is uniformly bounded away from zero, the set $\{\theta \geq \theta_\infty + \nu_0 : \frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x) \leq \eta\}$ must be empty for all t large enough, a contradiction.

This concludes our proof of part (i). For part (ii), let $(x_t^T)_{t\geq 0}$ be an optimal demand path for the game with horizon T, for each T. Let $P_{s,t}^T(\theta; x^T)$, $F_s^T(\theta; x^T)$, $V_t^T(\theta; x^T)$ be the analogous objects to $P_{s,t}(\theta; x)$, $F_s(\theta; x)$, $V_t(\theta; x)$ for the game with horizon T and demand path x^T. The proposition then states that $F_{T+1}^T(\overline{\theta}; x^T) \xrightarrow[T\to\infty]{} 0$. Suppose for the sake of contradiction that this is false, so that there is a subsequence $(T_l)_l$ and a fixed $\nu > 0$ such that $F_{T_l+1}^{T_l}(\overline{\theta}; x^{T_l}) \geq \nu$ for all l.

The same arguments as in part (i) imply that all optimal demands lie in the compact set $[0, \overline{x}]$. Then, by a diagonal argument and the Bolzano-Weierstrass theorem, we can find a subsequence $(T_{l_m})_m$ of $(T_l)_l$—denoted $(T_m)_m$ to save on notation—such that $F^{T_m}_{T_m+1}(\overline{\theta}; x^{T_m}) \geq \nu$ for all m and, for each $t \geq 0$, $(x^{T_m}_t)_m$ converges as $m \to \infty$ to a limit which we denote by x^∞_t. In addition, the condition $F^{T_m}_{T_m+1}(\overline{\theta}; x^{T_m}) \geq \nu$ for all m implies that there is some $\tilde{\theta} < \overline{\theta}$ and $\tilde{\nu} > 0$ such that $P^{T_m}_{T_m}(\theta; x^{T_m}) \geq \tilde{\nu}$ for all m and all $\theta \geq \tilde{\theta}$. (This relies on the fact that $P^T_t(\theta; x)$ is always weakly increasing in θ.)

For any fixed t and θ, it is clear that $V^{T_m}_{t+1}(\theta; x^{T_m}) \xrightarrow[m\to\infty]{} V_{t+1}(\theta; x^\infty)$ and thus $P^{T_m}_{t,t}(\theta; x^{T_m}) \xrightarrow[m\to\infty]{} P_{t,t}(\theta; x^\infty)$. Hence, for any fixed $s \geq t$ and θ, $P^{T_m}_{s,t}(\theta; x^{T_m}) \xrightarrow[m\to\infty]{} P_{s,t}(\theta; x^\infty)$. In particular, $P^{T_m}_s(\theta; x^{T_m}) \xrightarrow[m\to\infty]{} P_s(\theta; x^\infty)$ for all θ, s. Since $P^{T_m}_s(\theta; x^{T_m}) \geq P^{T_m}_{T_m}(\theta; x^{T_m})$ for all $s \leq T_m$, it follows that $P_s(\theta; x^\infty) \geq \tilde{\nu}$ for all $\theta \geq \tilde{\theta}$ and all s.

Since x^∞ does not induce receiver exit with probability 1, it follows from part (i) that it is a strictly suboptimal demand path in the game with infinite horizon. But, for all θ, $U^{T_m}_0(\theta; x^{T_m}) \xrightarrow[m\to\infty]{} U_0(\theta; x^\infty)$. By the dominated convergence theorem, the same holds for the proposer's utility unconditional on the receiver's type: $U^{T_m}_0(x^{T_m}) \xrightarrow[m\to\infty]{} U_0(x^\infty)$. Let x^* be an optimal demand path for $T = \infty$, so that $U_0(x^\infty) < U_0(x^*)$. Let x^{*T_m} be a truncation of x^* to horizon T_m. By the same arguments as above, $U^{T_m}_0(x^{*T_m}) \xrightarrow[m\to\infty]{} U_0(x^*)$, whence $U^{T_m}_0(x^{*T_m}) > U^{T_m}_0(x^{T_m})$ for some m, a contradiction. $\square$

Proof of Proposition 5. Consider first the case in which all demands by the proposer are interior, that is, $x_s > 0$ for all s.[6]

Under this assumption, as in Proposition 4, for each s, the solution must satisfy the FOC (9):

$$\pi'(x_s)F_{s+1}(\overline{\theta}; x) =$$
$$\sum_{t=0}^{s} \left(\frac{\delta_2}{\delta_1}\right)^{s-t} \int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{t-1}(\theta; x)g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right) P_{s,t+1}(\theta)\overline{U}_t(\theta; x)d\theta$$

Since $\delta_2 \geq \delta_1$, $\pi'(x_s) \leq \pi'(0)$, $f(\theta) \geq \underline{f}$, $P_{t-1}(\theta; x)P_{s,t+1}(\theta; x) \geq P_{t-1}(\theta; x)P_{s,t}(\theta; x) =$

[6] As argued in Proposition 4, any optimal demand must satisfy $x_s < \frac{\overline{\theta}+\eta}{1-\delta_2}$, but it is no longer obvious that positive demands are necessarily optimal in all periods: in particular, if $\delta_2 > \delta_1$, the proposer may now prefer to make low or even zero demands in late periods to "backload" incentives efficiently.

$P_s(\theta; x)$, and $\overline{U}_t(\theta; x) \geq \underline{M} := \frac{\theta - \eta}{1-\delta} > 0$, this implies that, for all s,

$$\pi'(0) F_{s+1}(\overline{\theta}; x) \geq \sum_{t=0}^{s} \int_{\underline{\theta}}^{\overline{\theta}} \underline{f} P_s(\theta; x) g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right) \underline{M} d\theta$$

Since $F_{s+1}(\overline{\theta}; x) = \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_s(\theta; x) d\theta \leq \overline{f} \int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) d\theta$, we have that, for all s,

$$\frac{\pi'(0)\overline{f}}{\underline{M}\underline{f}} \geq \frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) \left[\sum_{t=0}^{s} g\left(\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)\right)\right] d\theta}{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) d\theta} \tag{11}$$

For each s, let $[\theta_{*s}, \theta_s^*]$ be the set of marginal receiver types, that is, $\frac{\theta_{*s}}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta_{*s}; x) = -\eta$, $\frac{\theta_s^*}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta_s^*; x) = \eta$. Let $\theta_\infty^* = \limsup_{s\to\infty} \theta_s^*$, $\theta_{*\infty} = \limsup_{s\to\infty} \theta_{*s}$. In addition, let $\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x) = \epsilon_t^*(\theta; x)$. Recall that $\frac{\partial \epsilon_t^*(\theta; x)}{\partial \theta} \in \left(1, \frac{1}{1-\delta_2}\right)$.

We now split the proof into two cases. First, suppose that the receiver quits with probability 1 on the equilibrium path, but $P_s(\overline{\theta}; x) \searrow P_\infty(\overline{\theta}; x) > 0$, i.e., the maximal receiver type $\overline{\theta}$ does not quit with probability 1. (Note that this does not necessarily contradict the previous assumption; the receiver quitting with probability 1, unconditional on type, is equivalent to $P_s(\theta; x)$ converging pointwise to zero as $s \to \infty$ for $\theta < \overline{\theta}$, as well as equivalent to $\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) d\theta \searrow 0$ as $s \to \infty$.) Note that

$$\begin{aligned}
P_s(\theta; x) &= \prod_{t=0}^{s} P_{t,t}(\theta; x) = \prod_{t=0}^{s} G\left(\epsilon_t^*(\theta; x)\right) \\
\frac{\partial P_s(\theta; x)}{\partial \theta} &= \sum_{t=0}^{s} g(\epsilon_t^*(\theta; x)) \frac{\partial \epsilon_t^*(\theta; x)}{\partial \theta} \prod_{u \leq s, u \neq t} G\left(\epsilon_u^*(\theta; x)\right) \leq \frac{1}{1-\delta_2} \sum_{t=0}^{s} g(\epsilon_t^*(\theta; x))
\end{aligned}$$

Then we can bound the numerator of the right-hand side of (11) away from zero as follows:

$$\begin{aligned}
&\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) \left[\sum_{t=0}^{s} g\left(\epsilon_t^*(\theta; x)\right)\right] d\theta \geq (1-\delta_2) \int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) \frac{\partial P_s(\theta; x)}{\partial \theta} d\theta = \\
&= \frac{1-\delta_2}{2} P_s(\theta; x)^2 \Big|_{\underline{\theta}}^{\overline{\theta}} = \frac{1-\delta_2}{2} P_s(\overline{\theta}; x)^2 \geq \frac{1-\delta_2}{2} P_\infty(\overline{\theta}; x)^2.
\end{aligned}$$

This expression is uniformly bounded away from zero as $s \to \infty$, while $\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) d\theta$, the denominator of the right-hand side of (11), goes to zero as $s \to \infty$ by assumption. Then (11) must be violated for s large enough, a contradiction.

Second, suppose that the receiver quits with probability 1 on the equilibrium path, and $P_s(\overline{\theta}; x) \searrow 0$ as $s \to \infty$. (Intuitively, in this case, even receiver types close to the maximal type exit the game quickly.) Note that we must still have $P_s(\overline{\theta}; x) > 0$ for all finite s; otherwise the proposer would be making an offer which is rejected for sure on the equilibrium path, which we know is suboptimal. Rewrite the right-hand side of (11) as

$$\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) \left[\sum_{t=0}^{s} g\left(\epsilon_t^*(\theta; x)\right)\right] d\theta}{P_s(\overline{\theta}; x)} \frac{P_s(\overline{\theta}; x)}{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta; x) d\theta}.$$

We will show that $\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)\left[\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\right]d\theta}{P_s(\overline{\theta};x)}$ is bounded away from zero, and that $\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)d\theta}{P_s(\overline{\theta};x)}$ goes to zero, as $s \to \infty$, which again will imply that (11) is violated for s large enough.

Claim 1. $\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)d\theta}{P_s(\overline{\theta};x)} \xrightarrow[s\to\infty]{} 0.$

Proof. Suppose the claim does not hold, for the sake of contradiction. Equivalently, suppose that $\int_{\underline{\theta}}^{\overline{\theta}} \frac{P_s(\theta;x)}{P_s(\overline{\theta};x)} d\theta$ does not converge to zero as $s \to \infty$. Clearly we would have $\int_{\underline{\theta}}^{\overline{\theta}} \frac{P_s(\theta;x)}{P_s(\overline{\theta};x)} d\theta \xrightarrow[s\to\infty]{} 0$ if $\frac{P_s(\theta;x)}{P_s(\overline{\theta};x)}$ converged pointwise to 0 for $\theta < \overline{\theta}$ (note that $\frac{P_s(\theta;x)}{P_s(\overline{\theta};x)}|_{\theta=\overline{\theta}} = 1$ by construction, and $\frac{P_s(\theta;x)}{P_s(\overline{\theta};x)} \leq 1$ for all $\theta \leq \overline{\theta}$ since $P_s(\theta; x)$ is weakly increasing in θ). Then there must be $\theta_0 \in [\underline{\theta}, \overline{\theta})$ for which $\frac{P_s(\theta_0;x)}{P_s(\overline{\theta};x)}$ does not converge to zero as $s \to \infty$.

Since $\frac{P_s(\theta_0;x)}{P_s(\overline{\theta};x)} = \prod_{t=0}^{s} \frac{G(\epsilon_t^*(\theta_0;x))}{G\left(\epsilon_t^*(\overline{\theta};x)\right)}$ and $\frac{G(\epsilon_t^*(\theta_0;x))}{G\left(\epsilon_t^*(\overline{\theta};x)\right)} \leq 1$ for all t, $\frac{P_s(\theta_0;x)}{P_s(\overline{\theta};x)}$ must converge downward to a (positive) limit as $s \to \infty$. Using that $\frac{\partial \epsilon_t^*(\theta;x)}{\partial \theta} \in \left(1, \frac{1}{1-\delta_2}\right)$, we can bound

$$\frac{P_s(\theta_0; x)}{P_s(\overline{\theta}; x)} \leq \prod_{t=0}^{s} \frac{G\left(\epsilon_t^*(\overline{\theta}; x) - (\overline{\theta} - \theta_0)\right)}{G\left(\epsilon_t^*(\overline{\theta}; x)\right)}. \tag{12}$$

Now note that, for any fixed $\Delta > 0$, the ratio $\frac{G(x-\Delta)}{G(x)}$ must be uniformly bounded below 1 across all $x \in (-\eta, \eta]$. Indeed, $\frac{G(x-\Delta)}{G(x)}$ is a continuous function of x which converges to 0 as $x \to -\eta$ (in fact, it is identically zero for $x \in (-\eta, -\eta + \Delta]$); it is weakly below 1 because G is a c.d.f.; and it cannot take the value 1 for any $x \in (-\eta, \eta]$, since this would imply $G(x) = G(x - \Delta)$ and hence that $(x - \Delta, x) \cap \operatorname{supp} G = \emptyset$, a contradiction of A1($\eta$). Taking $\Delta = \overline{\theta} - \theta_0$, we obtain that the right-hand side of (12) goes to zero as $s \to \infty$, a contradiction, unless there are only finitely many values of t for which $\epsilon_t^*(\overline{\theta}; x) \leq \eta$. But, if so, then $P_s(\overline{\theta}; x)$ would not converge to zero as $s \to \infty$, a contradiction. □

Claim 2. $\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)\left[\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\right] d\theta}{P_s(\overline{\theta};x)}$ *is bounded away from zero for all s.*

Proof. Let $\hat{\theta} = \sup_s \theta_{*s}$. Recall that $\theta_{*s} \geq \overline{\theta}$ would imply an offer in period s that is rejected for sure, which must be suboptimal. (At any rate, the first such offer must be suboptimal.) Moreover, suppose that $\theta_{*\infty} = \overline{\theta}$. Then there are arbitrarily large times s for which the receiver's acceptance probability conditional on still being in the game in period s, $\frac{F_{s+1}(\overline{\theta};x)}{F_s(\overline{\theta};x)}$ is arbitrarily close to zero. Indeed, this is true along any sequence of times for which θ_{*s} converges to $\overline{\theta}$; it follows from A1(η) and the fact that the derivative of the expression $\frac{\theta}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta; x)$ with respect to θ is bounded by $\frac{1}{1-\delta_2}$. Since the proposer's maximum continuation payoff is bounded conditional on acceptance, this implies that the proposer's continuation utility, $\frac{\int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{s-1}(\theta;x) U_s(\theta;x) d\theta}{\int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{s-1}(\theta;x) d\theta}$, goes to zero along such a sequence. But the proposer can obtain a continuation payoff bounded away from zero by proposing $\underline{\theta} - \eta > 0$ in any such period, and this also improves the receiver's behavior in earlier periods, so it must be a profitable deviation. Hence $\theta_{*s} < \overline{\theta}$ for all s and $\theta_{*\infty} < \overline{\theta}$, so $\hat{\theta} < \overline{\theta}$. Then $\epsilon_t^*(\theta; x) \geq -\eta$ for all $\theta \geq \hat{\theta}$ and all t.

From the preceding argument, and the fact that $\frac{\partial \epsilon_t^*(\theta;x)}{\partial \theta} \in \left(1, \frac{1}{1-\delta_2}\right)$, it follows that $P_{t,t}(\theta; x) = G(\epsilon_t^*(\theta; x)) \geq G(-\eta + \theta - \hat{\theta})$ for all $\theta \geq \hat{\theta}$ and all t. Recall now that

$$\frac{\partial P_s(\theta; x)}{\partial \theta} = \sum_{t=0}^{s} g(\epsilon_t^*(\theta; x)) \frac{\partial \epsilon_t^*(\theta; x)}{\partial \theta} \prod_{u \leq s, u \neq t} G\left(\epsilon_u^*(\theta; x)\right) \leq$$

$$\leq \frac{1}{1-\delta_2} \sum_{t=0}^{s} g(\epsilon_t^*(\theta; x)) \prod_{u \leq s, u \neq t} G\left(\epsilon_u^*(\theta; x)\right).$$

For $\theta > \hat{\theta}$, we can further bound the above as follows:

$$\frac{1}{1-\delta_2}\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\prod_{u\leq s, u\neq t} G\left(\epsilon_u^*(\theta;x)\right) = \frac{1}{1-\delta_2}\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\frac{P_s(\theta;x)}{G(\epsilon_t^*(\theta;x))} \leq$$

$$\leq \frac{1}{1-\delta_2}\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\frac{P_s(\theta;x)}{G(-\eta+\theta-\hat{\theta})} = \frac{1}{1-\delta_2}\frac{1}{G(-\eta+\theta-\hat{\theta})}P_s(\theta;x)\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))$$

Let $\tilde{\theta} = \frac{\hat{\theta}+\overline{\theta}}{2}$ and $G_0 = G(-\eta+\tilde{\theta}-\hat{\theta}) > 0$. Then, for $\theta \geq \tilde{\theta}$,

$$\frac{\partial P_s(\theta;x)}{\partial \theta} \leq \frac{1}{(1-\delta_2)G_0}P_s(\theta;x)\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x)).$$

Then

$$\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)\left[\sum_{t=0}^{s} g\left(\epsilon_t^*(\theta;x)\right)\right] d\theta \geq \int_{\tilde{\theta}}^{\overline{\theta}} P_s(\theta;x)\left[\sum_{t=0}^{s} g\left(\epsilon_t^*(\theta;x)\right)\right] d\theta \geq$$

$$\geq (1-\delta_2)G_0\int_{\tilde{\theta}}^{\overline{\theta}} \frac{\partial P_s(\theta;x)}{\partial \theta} d\theta = (1-\delta_2)G_0(P_s(\overline{\theta};x) - P_s(\hat{\theta};x)).$$

Then $\frac{\int_{\underline{\theta}}^{\overline{\theta}} P_s(\theta;x)\left[\sum_{t=0}^{s} g(\epsilon_t^*(\theta;x))\right]d\theta}{P_s(\overline{\theta};x)} \geq (1-\delta_2)G_0\left[1 - \frac{P_s(\hat{\theta};x))}{P_s(\overline{\theta};x)}\right]$ for all s. This expression converges to $(1-\delta_2)G_0 > 0$ as $s \to \infty$, since $\frac{P_s(\hat{\theta};x)}{P_s(\overline{\theta};x)}$ goes to zero, as shown in the proof of Claim 1. $\square$

Let us now consider the alternative case in which the proposer's optimal demand path involves zero-demands, i.e., $x_s = 0$ for some s. If there is a finite s_0 such that $x_s = 0$ for all $s \geq s_0$, then of course the receiver never quits from period s_0 on. The result that the receiver does not exit with probability 1 follows immediately from the fact that it is never optimal for the proposer to make a demand that is rejected with probability 1: indeed, we must have $F_\infty(\overline{\theta};x) = F_{s_0}(\overline{\theta};x) > 0$.

The same argument applies more generally whenever there is a finite s_0 such that no receiver type close to $\overline{\theta}$ quits after s_0, i.e., if there is s_0 and $\hat{\theta} < \overline{\theta}$ such that $\theta_s^* \leq \hat{\theta}$ for all $s \geq s_0$. Suppose this is not the case, i.e., $\limsup \theta_s^* \geq \overline{\theta}$. Then we can split the set of periods $\mathbb{N}_0$ into two disjoint sets, A and B, where A contains all the periods in which the receiver quits with positive probability, and B contains the rest, i.e., $s \in A$ iff $F_{s+1}(\overline{\theta};x) < F_s(\overline{\theta};x)$; and the fact that $\limsup \theta_s^* \geq \overline{\theta}$ implies that A is infinite.

Moreover, since receivers never quit when flow payoffs are favorable, we must have $x_s \geq \underline{\theta} - \eta > 0$ for all $s \in A$, whence $\overline{U}_s(\theta; x) \geq \underline{\theta} - \eta > 0$ for all θ and all $s \in A$. In addition, since $x_s > 0$ for all $s \in A$, the FOC (9) must hold with equality for all $s \in A$. Then we can apply the same argument as in the main proof to this case, by exploiting the fact that (11) must hold for all $s \in A$.[7] Indeed, all the arguments that follow go through if (11) holds for an infinite number of periods rather than all of them. □

Proof of Proposition 6. We begin with an observation: it is without loss of generality to focus on direct revelation mechanisms where the receiver reveals θ and then the proposer imposes a demand path $(x_t(\theta))_{t\geq 0}$ satisfying a set of IC constraints. The reason is as follows. Given a general contract $X_t(\cdot)$, we can denote by $(x_t(\theta))_{t\geq 0}$ the path of demands that a receiver of type θ would opt into on the equilibrium path. Crucially, even though the receiver receives interim information about payoffs (ϵ_t is only observed in period t), and this information affects her decision to stay or quit, it *does not* affect her ranking of non-exit options; the realization of ϵ_t shifts the payoff obtained from all allowed transfers x_t equally, and leaves only the payoff from the exit option unchanged. Then the contract $\tilde{X}$ given by $\tilde{X}_0 = \{x_0(\theta) : \theta \in [\underline{\theta}, \overline{\theta}]$ and $\tilde{X}_t(x_t, \ldots, x_0) = \{x_t(\theta)\}$ if $x_0 = x_0(\theta)$ for $t > 0$ is outcome-equivalent to X.

As before, let $P_t(\theta)$ be θ's probability of accepting through period t, now in response to her personalized demand path. Let $P_{t,t}(\theta)$ be her probability of accepting *in* period t. Let $V(\tilde{\theta}; \theta)$ be the receiver's value function ex ante from reporting $\tilde{\theta}$ if her true type is θ. Then type θ's IC constraint is $V(\theta; \theta) \geq V(\tilde{\theta}; \theta)$ for all $\tilde{\theta}$. Denote $V(\theta; \theta) = V(\theta)$. Denote by $W(\tilde{\theta}; \theta)$ the utility of a type θ who reports to be type $\tilde{\theta}$ and then makes exit decisions as if she were type $\tilde{\theta}$ rather than optimally. (Of course, $W(\tilde{\theta}; \theta) \leq V(\tilde{\theta}; \theta)$.)

By the envelope theorem in integral form,[8] $V(\theta) \equiv \int_{\underline{\theta}}^{\theta} \frac{\partial V(\tilde{\theta};\tilde{\theta})}{\partial 2} d\tilde{\theta}$, where $\frac{\partial V(\theta;\theta)}{\partial 2} = -\sum_{t\geq 0} \delta_2^t (1 - P_t(\theta))$. Moreover, a contract $(x_t(\theta))_{t\geq 0, \theta \in [\underline{\theta}, \overline{\theta}]}$ is IC if and only if (i) $\frac{\partial V(\theta;\theta)}{\partial 2} = -\sum_{t\geq 0} \delta_2^t (1 - P_t(\theta))$ is a weakly increasing function of θ, which implies that $V(\theta)$ is weakly convex, and (ii) $V'(\theta^-) \leq \frac{\partial V(\theta;\theta)}{\partial 2} \leq V'(\theta^+)$ for all $\theta \in (\underline{\theta}, \overline{\theta})$, plus $\frac{\partial V(\underline{\theta};\underline{\theta})}{\partial 2} \leq V'(\underline{\theta}^+)$ and $V'(\overline{\theta}^-) \leq \frac{\partial V(\overline{\theta};\overline{\theta})}{\partial 2}$. To see why, note that, if $-\sum_{t\geq 0} \delta_2^t (1 - P_t(\theta)) >$

[7] Note that the expression $\overline{U}_t$ shows up in the right-hand side of (9) for all t, not just $t \in A$, even if $s \in A$. However, for $t \in B$, the corresponding term in the right-hand side of (9) is zero since no remaining receiver types are marginal in that period, so we can bound $\overline{U}_t \geq \underline{\theta} - \eta > 0$ for all relevant t.

[8] See Theorem 2 in Milgrom and Segal (2002).

$-\sum_{t\geq 0}\delta_2^t(1-P_t(\theta'))$ for some $\theta<\theta'$, then

$$W(\theta';\theta)+W(\theta;\theta')=V(\theta;\theta)+V(\theta';\theta')+(\theta'-\theta)(\sum_{t\geq 0}\delta_2^t(P_t(\theta)-P_t(\theta'))>V(\theta';\theta)+V(\theta;\theta'),$$

which implies that either $V(\theta';\theta)\geq W(\theta';\theta)>V(\theta;\theta)$ or $V(\theta;\theta')\geq W(\theta;\theta')>V(\theta';\theta')$, so one of these types can profitably deviate. Similarly, if $\frac{\partial V(\theta;\theta)}{\partial 2}>V'(\theta^+)$ for some θ, then types in a right-neighborhood of θ can profit by reporting type θ, and analogously if $V'(\theta^-)>\frac{\partial V(\theta;\theta)}{\partial 2}$. [[[Deal with sufficiency]]]

The demand path offered to a type θ under the optimal contract must then also be a solution to the following constrained problem: given a type θ and values V_0 and D_0, what path $(x_t(\theta))_{t\geq 0}$ maximizes the principal's objective

$$\sum_{t=0}^{\infty}\delta_1^t\pi(x_t(\theta))P_t(\theta),$$

subject to the constraints that θ's utility when facing the demand path $(x_t(\theta))_{t\geq 0}$ is V_0, and $\sum_{t\geq 0}\delta_2^t P_t(\theta)=D_0$?

The Lagrangian for this problem is:

$$\mathcal{L}=\sum_{t=0}^{\infty}\delta_1^t\pi(x_t(\theta))P_t(\theta)+\lambda(V_0-V(\theta))+\mu(D_0-\sum_{t=0}^{\infty}\delta_2^tP_t(\theta))$$

where $V(\theta)$ is θ's equilibrium utility given the demand path $x(\theta)$.

Consider a deviation of the following form: x_{t-1} changes by $-\nu\delta_2P_{t,t}(\theta)$, and x_t changes by ν, for $\nu>0$ small. The impact of this deviation on the receiver's continuation value $V_s(\theta)$ in each period s, conditional on still being in the game at the beginning of that period, is as follows. For $s>t$, V_s is clearly unchanged. For $s=t$, $\frac{\partial V_t}{\partial\nu}=-P_{t,t}(\theta)$ by the envelope theorem. For $s=t-1$, $\frac{\partial(-x_{t-1}+\delta_2V_t)}{\partial\nu}|_{\nu=0}=\delta_2P_{t,t}(\theta)-\delta_2P_{t,t}(\theta)=0$, so $\frac{\partial V_{t-1}}{\partial\nu}|_{\nu=0}=0$. Then $\frac{\partial V_s}{\partial\nu}|_{\nu=0}=0$ for all $s<t$. In particular, $\frac{\partial V_0}{\partial\nu}|_{\nu=0}=0$.

The impact on this deviation on $P_{s,s}(\theta)$ in each period is as follows. For $s>t$, $P_{s,s}(\theta)$ is unchanged. For $s=t$, $\frac{\partial P_{t,t}(\theta)}{\partial\nu}|_{\nu=0}=-g\left(\frac{\theta}{1-\delta_2}-x_t+\delta_2V_{t+1}\right)=:-g_t$ by (7). For $s<t$, $\frac{\partial P_{s,s}(\theta)}{\partial\nu}|_{\nu=0}=0$ because there is no change to $-x_{t-1}+\delta_2V_t$ and no change for x_{s-1}, V_s for $s<t$.

The impact on $P_s(\theta)$ for each s is then as follows. For $s<t$, $\frac{\partial P_s}{\partial\nu}|_{\nu=0}=0$.

For $s = t$, $\frac{\partial P_t(\theta)}{\partial \nu}|_{\nu=0} = P_{t-1}(\theta)\frac{\partial P_{t,t}(\theta)}{\partial \nu}|_{\nu=0} = -P_{t-1}(\theta)g_t = -\frac{g_t}{P_{t,t}(\theta)}P_t(\theta)$. For $s > t$, $\frac{\partial P_s(\theta)}{\partial \nu}|_{\nu=0} = P_s(\theta)\frac{\frac{\partial P_t(\theta)}{\partial \nu}|_{\nu=0}}{P_t(\theta)} = -\frac{g_t}{P_{t,t}(\theta)}P_s(\theta)$.

Denote $\pi(x_s) = \pi_s$, $\pi'(x_s) = \pi'_s$, $P_s(\theta) = P_s$, $P_{s,s}(\theta) = P_{s,s}$. Then

$$\frac{\partial \mathcal{L}}{\partial \nu}|_{\nu=0} = -\delta_1^{t-1}\pi'_{t-1}\delta_2 P_{t,t}P_{t-1} + \delta_1^t \pi'_t P_t + \sum_{s\geq t}\delta_1^s \pi_s \frac{\partial P_s}{\partial \nu}|_{\nu=0} - \mu\sum_{s\geq t}\delta_2^s \frac{\partial P_s}{\partial \nu}|_{\nu=0}$$

$$\implies \frac{1}{\delta_1^t}\frac{\partial \mathcal{L}}{\partial \nu}|_{\nu=0} = (-\frac{\delta_2}{\delta_1}\pi'_{t-1} + \pi'_t)P_t + \frac{g_t}{P_{t,t}}\left(-\sum_{s\geq t}\delta_1^{s-t}\pi_s P_s + \mu\sum_{s\geq t}\left(\frac{\delta_2}{\delta_1}\right)^t \delta_2^{s-t}P_s\right).$$

If the allocation is optimal, the right-hand side should vanish for all t.[9]

Suppose type θ stays forever with positive probability, so $P_s \searrow P_\infty > 0$. Then $P_{t,t} \to 1$. Now, $\mu\sum_{s\geq t}\left(\frac{\delta_2}{\delta_1}\right)^t \delta_2^{s-t}P_s$ is bounded above by $\mu\left(\frac{\delta_2}{\delta_1}\right)^t \frac{1}{1-\delta_2} \xrightarrow[t\to\infty]{} 0$. On the other hand, $\sum_{s\geq t}\delta_1^{s-t}\pi_s P_s$ is bounded above by $\frac{\overline{\theta}+\eta}{1-\delta_2}\frac{1}{1-\delta_1}P_s$, so the limsup of the second term in absolute value is no more than $\frac{\overline{\theta}+\eta}{1-\delta_2}\frac{1}{1-\delta_1}P_\infty \limsup_t g_t$.

But, if $P_{t,t} \to 1$, then $g_t \to 0$ by A2(η), so this expression goes to zero. Then $-\frac{\delta_2}{\delta_1}\pi'_{t-1} + \pi'_t \to 0$, which implies that $\pi'_t \to 0$. This is either impossible or leads to $x_t \to \infty$, contradicting $P_\infty > 0$.

□

Proof of Proposition 8. Part (i): we proceed again by contradiction. Borrowing notation from Proposition 4, suppose for the sake of contradiction that the receiver stays in the game forever with positive probability, so $F_\infty(\overline{\theta}; x) > 0$, and let $\theta_\infty = \inf(\text{supp}(P_\infty(\theta; x)))$. As argued in Proposition 4, we must have $\liminf_{t\to\infty} \frac{\theta_\infty}{1-\delta_2} - x_t + \delta_2 V_{t+1}(\theta_\infty; x) \geq \eta$. Moreover, a proposal $x_s > \overline{x}$ is always rejected and hence suboptimal, while a proposal $x_s = 0$ is suboptimal because any proposal $x_s \in (0, \underline{\theta} - \eta)$ is always accepted.[10]

As a result, equilibrium proposals must always be interior, and hence satisfy, for

[9] Technically the proposer's problem includes the additional constraints $x_s(\theta) \geq 0$ for all s. However, it is easy to show that, if $x_s(\theta) < \underline{\theta}$, then the proposer can always increase her payoff by increasing $x_s(\theta)$ by a small $\nu > 0$ and decreasing $x_{s-1}(\theta)$ by $\nu P_{s,s}(\theta)$, [[[]]]

[10] Note that even deviating from $x_s = 0$ to $x_s > 0$ could in principle be undesirable if the receiver thought such a deviation presaged further escalation; the assumption of a naive receiver rules out this possibility.

each $s \geq 0$, the FOC:

$$\pi'(x_s)F_{s+1}(\overline{\theta};x) = \int_{\underline{\theta}}^{\overline{\theta}} f(\theta)P_{s-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta;x)\right)\overline{U}_s(\theta;x)d\theta. \tag{13}$$

Note that this FOC is identical to equation (9), except that the only term on the right-hand side is the one corresponding to $t = s$. In other words, the only cost of increasing x_s the proposer takes into account is that a higher x_s increases the receiver's probability of exit in period s; a deviation at time s would have no effect in earlier periods as it would be unanticipated by the receiver (and, by time s, the receiver's behavior before period s is sunk). At the same time, the assumption that the receiver is naive guarantees that a deviation at time s has no impact on the receiver's continuation value V_{s+1} or her expected behavior in periods $s+1$ and onwards.

The same approach as in Proposition 4 then yields a contradiction—in this case, regardless of the values of δ_1 and δ_2. Indeed, (13) implies that, for all s,

$$\pi'(\overline{x})F_\infty(\overline{\theta};x) \leq \overline{f}M\int_{\underline{\theta}}^{\overline{\theta}} P_{s-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta;x)\right)d\theta.$$

Since $\int_{\underline{\theta}}^{\theta_\infty} P_{s-1}(\theta;x) \to 0$ and $g(\cdot) \leq \overline{g}$, this implies

$$\begin{aligned} &\frac{\pi'(\overline{x})F_\infty(\overline{\theta};x)}{\overline{f}M} \leq \liminf_{s\to\infty}\int_{\theta_\infty}^{\overline{\theta}} P_{s-1}(\theta;x)g\left(\frac{\theta}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta;x)\right)d\theta \\ \implies &\frac{\pi'(\overline{x})F_\infty(\overline{\theta};x)}{\overline{f}M} \leq \liminf_{s\to\infty}\int_{\theta_\infty}^{\overline{\theta}} g\left(\frac{\theta}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta;x)\right)d\theta \end{aligned}$$

By the same argument as in Proposition 4, the right-hand side converges to zero, a contradiction.

Part (ii): As argued before, proposals $x_s \geq \overline{x}$ or $x_s = 0$ are always strictly suboptimal.[11] Suppose for the sake of contradiction that $F_\infty(\overline{\theta};x) > 0$.

[11] In this case, the argument for the suboptimality of $x_s = 0$ is slightly different. Consider a deviation from $x_s = 0$ to $x_s = \nu$, for $\nu > 0$ arbitrarily small. With probability $1 - G(\eta - \nu)$, the proposer gains ν at no cost. With probability $G(\eta - \nu)$, a different outcome may materialize; the worst that can happen is that whenever $y_s > \eta$ the receiver quits. Even then, the cost of this deviation is at most $MG(\eta - \nu)$. By A2(η), $M\int_{\eta-\nu}^{\eta} g(x)dx < \nu(1 - \int_{\eta-\nu}^{\eta} g(x)dx)$ for ν small enough, since g is continuous and $g(\eta) = 0$.

For each $s \geq 0$, a deviation from x_s to $x_s + \nu$ must be weakly unprofitable for any $\nu > 0$. Note that, whenever $y_t \in [x_t - \nu, x_t + \nu]$ for all $t \leq s$, the receiver's belief at time s is that with probability 1 the proposer has not deviated and will not deviate in the future, so her continuation value from acceptance at time s is $\frac{\theta}{1-\delta_2} - y_s + \delta_2 V_{s+1}(\theta; x)$. If $y_t \notin [x_t - \nu, x_t + \nu]$ for some $t \leq s$, the receiver's belief is not pinned down by Bayes' rule.

Suppose that whenever $y_t \notin [x_t - \nu, x_t + \nu]$, the receiver quits immediately in period t. This extreme response may not be incentive-compatible, but it is a worst-case scenario for the proposer. We will show that even if the receiver responses to detecting a deviation in this way, the proposer will be tempted to deviate whenever $F_\infty(\overline{\theta}; x) > 0$. Suppose instead that the proposer does not want to increase x_s for any s. Then we must have that, for all s,

$$\pi'(x_s) F_{s+1}(\overline{\theta}; x) \leq \int_{\underline{\theta}}^{\overline{\theta}} f(\theta) P_{s-1}(\theta; x) g\left(\frac{\theta}{1-\delta_2} - x_s + \delta_2 V_{s+1}(\theta; x)\right) \overline{U}_s(\theta; x) d\theta.$$

Note that the possibility of a marginal deviation being revealed to the receiver (i.e., a realization $y_s > x_s + \eta$) does not appear in this FOC. The reason is that, for a deviation of size $\nu > 0$ with ν arbitrarily small, the probability that a deviation is revealed is $1 - G(\eta - \nu) \in o(\nu)$, by A2($\eta$). The cost to the proposer if this event does come to pass is at most M. The same argument as in part (i) can then be applied to obtain a contradiction, since the FOCs are essentially identical.

$\square$

Proof of Proposition 9. Consider an equilibrium defined by functions x^*, θ^* as described in the text. Suppose that the receiver does not quit with probability 1, that is, there is positive probability that $\theta^*_\infty = \lim_{t\to\infty} \theta^*_t < \overline{\theta}$.

We first show a contradiction in the following case in which the argument is most transparent. Suppose that, on the equilibrium path, the state never increases from a certain value θ, that is, $\theta^*(y) \leq \theta$ for all $y \in [x^*(\theta) - \eta, x^*(\theta) + \eta]$. Suppose $\theta_t = \theta$ and the proposer considers a deviation to setting $x_t = x^*(\theta) + \nu$ for a small $\nu > 0$. Then, with probability $G(\eta - \nu)$, $\epsilon_t \leq \eta - \nu$, whence $y_t \in [x^*(\theta) - \eta + \nu, x^*(\theta) + \eta]$ and no receiver type quits, so that $\theta_{t+1} = \theta$, as it would have been without a deviation. For small ν, this deviation yields a gain of at least $\pi'(\overline{x}) G(\eta - \nu)\nu$ and a cost of at most $(1 - G(\eta - \nu))M$. It then follows from A2(η) that this deviation is profitable for $\nu > 0$ small, since $G(\eta - \nu) \in o(\nu)$.

More generally, this argument applies if $\theta^*(y) \leq \theta$ for all $y \in [x^*(\theta) - \eta, x^*(\theta) + \eta)$. Suppose next that there is $\hat{\theta} < \overline{\theta}$ such that, letting $\hat{x}^- = \lim_{\theta \nearrow \hat{\theta}} x^*(\theta)$, we have $\lim_{\epsilon \nearrow \eta} \theta^*(\hat{x}^- + \epsilon) \leq \hat{\theta}$. Clearly, if $\theta_t < \hat{\theta}$, then the state remains weakly below $\hat{\theta}$ forever with probability 1. But this again leads to a contradiction. Indeed, suppose that the state is $\theta < \hat{\theta}$ and the proposer deviates from $x^*(\theta)$ to $\hat{x}^- + \nu$ for some $\nu > 0$. With probability $\frac{1-F(\hat{\theta})}{1-F(\theta)} G(\eta - \nu)$, the receiver's true type is at least $\hat{\theta}$, and the shock realization is less than $\eta - \nu$, whence the proposer's effective demand is less than $\hat{x}^- + \eta$, and the receiver accepts. The proposer thus gains at least $\pi'(\overline{x}) \frac{1-F(\hat{\theta})}{1-F(\theta)} G(\eta - \nu)\nu$ from this deviation in the current period. The cost is not more than $\left(\frac{F(\hat{\theta}) - F(\theta)}{1-F(\theta)} + (1 - G(\eta - \nu)) \right) M$—we obtain this upper bound by assuming that, whenever the receiver's type is below $\hat{\theta}$ *or* the shock realization is at least $\eta - \nu$, the receiver rejects at once. Note also that, with some probability, this deviation may lead to the state moving from θ to some $\theta' \in (\theta, \hat{\theta}]$, without the receiver rejecting. This weakly increases the proposer's payoff, since he can at worst make the same offers he would have made if the state did not increase and they will be accepted with weakly higher probability, so we can ignore it.

Take ν small enough that $\pi'(\overline{x})\nu > 2(1 - G(\eta - \nu))M$. Then we obtain a contradiction for θ close enough to $\hat{\theta}$.

For each θ, let $x^-(\theta) = \lim_{\theta' \nearrow \theta} x^*(\theta')$, and let $S(\theta) = \lim_{\epsilon \nearrow \eta} \theta^*(x^-(\theta) + \epsilon)$. The preceding argument implies that $S(\theta) > \theta$ for all $\theta < \overline{\theta}$. Note also that S is weakly increasing, by our assumption that x^* and θ^* are weakly increasing. Moreover, suppose that, for some θ, $S^k(\theta) \nearrow \hat{\theta} < \overline{\theta}$, where S^k denotes the composition of S with itself k times. We will argue that this implies $S(\hat{\theta}) = \hat{\theta}$ (leading to a contradiction). Indeed, we can rewrite

$$S(\hat{\theta}) = \sup_{\theta' < \hat{\theta}, \epsilon < \eta} \theta^*(x^*(\theta') + \epsilon) = \sup_k \left(\sup_{\theta' < S^{k-1}(\theta), \epsilon < \eta} \theta^*(x^*(\theta') + \epsilon) \right) = \sup_k S^k(\theta) = \hat{\theta}.$$

Thus, for any $\theta \leq \overline{\theta}$, $\lim_k S^k(\theta) \geq \overline{\theta}$.

Next, we will argue that, for any $\hat{\theta} < \overline{\theta}$, there is $\rho > 0$ such that $S(\theta) \geq \theta + \rho$ for all $\theta \in [\underline{\theta}, \hat{\theta}]$. Suppose not, so there is a sequence $(\theta_k)_k \subseteq [\underline{\theta}, \hat{\theta}]$ such that $S(\theta_k) - \theta_k \to 0$. By Bolzano-Weierstrass, there is a convergent subsequence $(\theta_{k_m})_m$, with a limit $\tilde{\theta} \in [\underline{\theta}, \hat{\theta}]$. Moreover, we can pick this subsequence to be strictly monotonic. (Indeed, we must have either $\theta_{k_m} > \tilde{\theta}$ for infinitely many values of m, or $\theta_{k_m} < \tilde{\theta}$ for infinitely many values of m, if not both. In the former case, we can select a subsequence converging to $\tilde{\theta}$ from above, and in the latter from below.)

If $\theta_{k_m} \nearrow \tilde{\theta}$, then $S(\theta_{k_m}) - \theta_{k_m} \to 0$ implies $S(\theta_{k_m}) \to \tilde{\theta}$, hence $S(\theta_{k_m}) \nearrow \tilde{\theta}$ by the monotonicity of S and $(\theta_{k_m})_m$. In turn, by the monotonicity of S and the fact that $\hat{\theta} = \sup_m \theta k_m$, this implies that $S(\theta) \leq \tilde{\theta}$ for all $\theta \leq \tilde{\theta}$, so $\lim_k S^k(\theta) \leq \tilde{\theta}$ for any $\theta < \tilde{\theta}$, a contradiction. If $\theta_{k_m} \searrow \tilde{\theta}$, then $S(\theta_{k_m}) - \theta_{k_m} \to 0$ implies $S(\theta_{k_m}) \to \tilde{\theta}$, hence $S(\theta_{k_m}) \searrow \tilde{\theta}$ by the monotonicity of S and $(\theta_{k_m})_m$. Then $S(\tilde{\theta}) \leq \tilde{\theta}$ by the monotonicity of S, a contradiction.

Let $Z(\theta) = \sup_{\epsilon<\eta} \theta^*(x^*(\theta) + \epsilon)$. By construction, Z is weakly increasing; $Z(\theta) \geq S(\theta)$ for all θ; and, as a result, $Z(\theta) \geq \theta + \rho$ for all $\theta \in [\underline{\theta}, \hat{\theta}]$. We will now show that, for any $\theta \in [\underline{\theta}, \hat{\theta}]$, if $\theta_t = \theta$, there is a positive probability that $\theta_{t+1} \geq \theta + \frac{\rho}{2}$. Indeed, since $\sup_{\epsilon<\eta} \theta^*(x^*(\theta) + \epsilon) \geq \theta + \rho$, there is $\epsilon_0 < \eta$ such that $\theta^*(x^*(\theta) + \epsilon_0) \geq \theta + \frac{\rho}{2}$, so $\theta_{t+1} \geq \theta + \frac{\rho}{2}$ with probability at least $1 - G(\eta - \epsilon_0) > 0$. As a result, the state *eventually* goes above $\theta + \frac{\rho}{2}$ with probability 1. (Note that if the state increases to a value between θ and $\theta + \frac{\rho}{2}$, x^* weakly increases, so the probability of the state going above $\theta + \frac{\rho}{2}$ in any future period also weakly increases.) Iterating this argument finitely many times yields that the state eventually goes above $\hat{\theta}$ with probability 1. Since this argument works for any $\hat{\theta} < \overline{\theta}$, all types $\theta < \overline{\theta}$ eventually quit with probability 1, as we wanted to show.

Derivation of Equation (3). Here is an informal argument. [[[Needs cleanup!]]] An optimal demand path must be such that the proposer demands $x_t = x^*(\theta_t^*)$ if the current marginal receiver type is θ_t^*, for some function $x^*(\cdot)$. Let $U(\theta)$ be the proposer's value function conditional on a marginal type θ. Then U solves the Bellman equation

$$U(\theta) = \max_x \left[G(\theta - x)(\pi(x) + \delta_1 U(\theta)) + \int_{\theta-x}^{\infty} g(\epsilon)(\pi(x) + \delta_1 U(x + \epsilon))(1 - F_\theta(x + \epsilon)) d\epsilon \right],$$

where $F_\theta(\tilde\theta) = \frac{F(\tilde\theta)-F(\theta)}{1-F(\theta)}\mathbb{1}_{\tilde\theta\geq\theta}$. Note that, by the envelope theorem,

$$
\begin{aligned}
U'(\theta) &= G(\theta-x)\delta_1 U'(\theta) + \int_{\theta-x}^{\infty} g(\epsilon)(\pi(x)+\delta_1 U(x+\epsilon))\frac{f(\theta)}{(1-F(\theta))^2}(1-F(x+\epsilon))d\epsilon \\
&= \frac{1}{1-\delta_1 G(\theta-x)}\frac{f(\theta)}{(1-F(\theta))^2}\int_{\theta-x}^{\infty} g(\epsilon)(\pi(x)+\delta_1 U(x+\epsilon))(1-F(x+\epsilon))d\epsilon \\
&= \frac{1}{1-\delta_1 G(\theta-x)}\frac{f(\theta)}{1-F(\theta)}\int_{\theta-x}^{\infty} g(\epsilon)(\pi(x)+\delta_1 U(x+\epsilon))(1-F_\theta(x+\epsilon))d\epsilon \\
&= \frac{1}{1-\delta_1 G(\theta-x)}\frac{f(\theta)}{1-F(\theta)}\left[U(\theta)-G(\theta-x)(\pi(x)+\delta_1 U(\theta))\right]
\end{aligned}
$$

If x is optimal, it must solve the FOC:

$$
\begin{aligned}
&0 = \pi'(x)G(\theta-x) + \int_{\theta-x}^{\infty} g(\epsilon)\left[(\pi'(x)+\delta_1 U'(x+\epsilon))(1-F_\theta(x+\epsilon)) - (\pi(x)+\delta_1 U(x+\epsilon))f_\theta(x+\epsilon)\right]d\epsilon \\
&\pi'(x)G(\theta-x) + \int_{\theta-x}^{\infty} g(\epsilon)(\pi'(x)+\delta_1 U'(x+\epsilon))(1-F_\theta(x+\epsilon))d\epsilon = \int_{\theta-x}^{\infty} g(\epsilon)(\pi(x)+\delta_1 U(x+\epsilon))f_\theta(x+\epsilon
\end{aligned}
$$

Suppose $x\in(\theta-\eta,\theta+\eta)$ and η is "small". Then

$$
\begin{aligned}
\pi'(\theta)G(\theta-x) + (\pi'(\theta)+\delta_1 U'(\theta))(1-G(\theta-x)) &= (\pi(\theta)+\delta_1 U(\theta))\frac{f(\theta)}{1-F(\theta)}\int_{\theta-x}^{\infty} g(\epsilon)d\epsilon \\
\pi'(\theta)G(\theta-x) + (\pi'(\theta)+\delta_1 U'(\theta))(1-G(\theta-x)) &= (\pi(\theta)+\delta_1 U(\theta))\frac{f(\theta)}{1-F(\theta)}(1-G(\theta-x)) \\
\pi' G + (\pi'+\delta_1 U')(1-G) &= (\pi+\delta_1 U)\frac{f}{1-F}(1-G) \\
\pi' + (1-G)\delta_1\frac{1}{1-\delta_1 G}\frac{f}{1-F}\left[U-(\pi+\delta_1 U)G\right] &= (\pi+\delta_1 U)\frac{f}{1-F}(1-G) \\
\frac{\pi'}{1-G} + \delta_1\frac{1}{1-\delta_1 G}\frac{f}{1-F}\left[U-(\pi+\delta_1 U)G\right] &= (\pi+\delta_1 U)\frac{f}{1-F} \\
\frac{\pi'}{1-G} + \delta_1\frac{f}{1-F}U - \delta_1\frac{1}{1-\delta_1 G}\frac{f}{1-F}\pi G &= \frac{f}{1-F}\pi + \delta_1\frac{f}{1-F}U \\
\frac{\pi'}{1-G} - \delta_1\frac{1}{1-\delta_1 G}\frac{f}{1-F}\pi G &= \frac{f}{1-F}\pi
\end{aligned}
$$

$$\begin{aligned}
\frac{\pi'}{1-G} &= \frac{f}{1-F}\pi\frac{1}{1-\delta_1 G} \\
\frac{1-\delta_1 G}{1-G} &= \frac{f}{1-F}\frac{\pi}{\pi'} \\
G &= \frac{\frac{f}{1-F}\frac{\pi}{\pi'}-1}{\frac{f}{1-F}\frac{\pi}{\pi'}-\delta_1} \\
G(\theta-x) &= \frac{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-1}{\frac{f(\theta)}{1-F(\theta)}\frac{\pi(\theta)}{\pi'(\theta)}-\delta_1}.
\end{aligned}$$

□

□

Derivation of x_0 and ω.

We can pin down x_0 and ω as follows. Let $z_0 = x_0(1-\omega)$. Let $V(\theta)$ be the value function of a receiver of type θ when she is marginal, i.e., when she is the lowest type left in the distribution. Then she must be indifferent when $\theta^*(y) = \theta$, i.e., when $\theta = y(1-\omega)$ or $y = \frac{\theta}{1-\omega}$, so

$$\begin{aligned}
-\frac{\theta}{1-\delta_2} &= -y + \delta_2 V(\theta) \\
V(\theta) &= \frac{1}{\delta_2}\theta\frac{\omega}{1-\omega} - \frac{\theta}{1-\delta_2}.
\end{aligned}$$

Let ϵ^* be the realization of ϵ for which $y = \frac{\theta}{1-\omega}$, i.e., $x_0\theta(1+\epsilon^*)(1-\omega) = \theta$, so

$\epsilon^* = \frac{1}{z_0} - 1$. Then $V(\theta)$ must satisfy the Bellman equation:

$$
\begin{aligned}
V(\theta) &= \int_{-\eta}^{\epsilon^*} \left[-x_0\theta(1+\epsilon) + \delta_2 V(\theta)\right] g(\epsilon)d\epsilon + \int_{\epsilon^*}^{\eta} -\frac{\theta}{1-\delta_2} g(\epsilon)d\epsilon \\
&= -x_0\theta \int_{-\eta}^{\epsilon^*} g(\epsilon)(1+\epsilon)d\epsilon + \delta_2 V(\theta)G(\epsilon^*) - \frac{\theta}{1-\delta_2}(1-G(\epsilon^*)) \\
(1-\delta_2 G(\epsilon^*))\left(\frac{1}{\delta_2}\theta\frac{\omega}{1-\omega} - \frac{\theta}{1-\delta_2}\right) &= -x_0\theta \int_{-\eta}^{\epsilon^*} g(\epsilon)(1+\epsilon)d\epsilon - \frac{\theta}{1-\delta_2}(1-G(\epsilon^*)) \\
(1-\delta_2 G(\epsilon^*))\frac{1}{\delta_2}\frac{\omega}{1-\omega} &= -x_0 \int_{-\eta}^{\epsilon^*} g(\epsilon)(1+\epsilon)d\epsilon + G(\epsilon^*) \\
(1-\delta_2 G(\epsilon^*))\frac{1}{\delta_2}\omega &= -z_0 \int_{-\eta}^{\epsilon^*} g(\epsilon)(1+\epsilon)d\epsilon + G(\epsilon^*)(1-\omega) \\
\frac{\omega}{\delta_2} &= -z_0 \int_{-\eta}^{\epsilon^*} g(\epsilon)(1+\epsilon)d\epsilon + G(\epsilon^*)
\end{aligned}
$$

On the other hand, note that the proposer's value function when the marginal type is θ, $U(\theta)$, must satisfy $U(\theta) \equiv \theta U_0$, where

$$
U_0 = x_0(G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon) + \delta_1 U_0 \left[G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right].
$$

This equation pins down U_0 given a value of x_0 and given that $z_0 = x_0(1-\omega)$, $\epsilon^* = \frac{1}{z_0} - 1$. Since x_0 must be optimal in equilibrium, we must have $\frac{dU_0}{dx_0} = 0$. (Note that varying x_0 amounts to having the proposer switch to a strategy demanding $x\theta$ when the marginal type is θ, for some $x \neq x_0$, no matter what θ is; but this deviation is profitable iff a one-shot deviation changing only the initial demand in this fashion is profitable.) Rewrite the previous equation as

$$
U_0 = \frac{x_0(G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon)}{1-\delta_1\left[G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right]}.
$$

Let $A = x_0(G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon)$, $B = 1-\delta_1\left[G(\epsilon^*) + \int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right]$. Then we must have $\frac{\frac{\partial A}{\partial x_0}}{A} = \frac{\frac{\partial B}{\partial x_0}}{B}$, i.e.,

$$\frac{G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon)-x_0\int_{\epsilon^*}^{\eta} g(\epsilon)f(z_0(1+\epsilon))(1-\omega)(1+\epsilon)d\epsilon}{x_0(G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon)}=$$

$$\frac{\delta_1\int_{\epsilon^*}^{\eta} g(\epsilon)f(z_0(1+\epsilon))z_0(1-\omega)(1+\epsilon)^2d\epsilon-\delta_1\int_{\epsilon^*}^{\eta} g(\epsilon)(1-\omega)(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon}{1-\delta_1\left[G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right]}$$

$$\frac{G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon)-z_0\int_{\epsilon^*}^{\eta} g(\epsilon)f(z_0(1+\epsilon))(1+\epsilon)d\epsilon}{G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon}=$$

$$\frac{\delta_1\int_{\epsilon^*}^{\eta} g(\epsilon)f(z_0(1+\epsilon))z_0^2(1+\epsilon)^2d\epsilon-\delta_1\int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon}{1-\delta_1\left[G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)z_0(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right]}$$

$$1-z_0\frac{\int_{\epsilon^*}^{\eta} g(\epsilon)(1+\epsilon)f(z_0(1+\epsilon))d\epsilon}{G(\epsilon^*)+\int_{\epsilon^*}^{\eta} g(\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon}=$$

$$\delta_1 z_0\frac{z_0\int_{\epsilon^*}^{\eta} g(\epsilon)(1+\epsilon)^2f(z_0(1+\epsilon))d\epsilon-\int_{\epsilon^*}^{\eta} g(\epsilon)(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon}{1-\delta_1\left[G(\epsilon^*)+z_0\int_{\epsilon^*}^{\eta} g(\epsilon)(1+\epsilon)(1-F(z_0(1+\epsilon)))d\epsilon\right]}$$

Proof of Proposition 10. A path of play for the receiver is a sequence $S=(p_t)_t$ of one of three possible types. We say $(p_t)_t$ is *type 1* if it is infinite, with $p_t\in\{p^0,\dots,p^{k-1}\}$ for all t, and only finitely many elements are nonzero. It is *type 2* if $p_t\in\{p^0,\dots,p^{k-1}\}$ for all t, but infinitely many elements are positive. It is *type 3* if it is finite, with (only) the last term equal to 1.

Let $P(\theta,S)$ be the probability that a receiver of type θ plays according to S on path (assuming that all dice rolls lead to peace). Let $P(S)=\int P(\theta,S)dF(\theta)$. Then we need to show for each part of the result that $P(S)=0$ for all type 1 sequences S under the conditions given.

(i) Suppose otherwise, and let S_0 be a sequence with minimal number of nonzero elements among all type 1 sequences S with $P(S)>0$. Denote by $S|t_0$ a sequence truncated to length t_0, i.e., if $S=(p_t)_{t\in\mathbb{N}_0}$ then $S|t_0=(p_0,\dots,p_{t_0-1})$.

Note that a proposer strategy can be described by a function $x_t(S^t)$ defining a demand x for all t and all length-t sequences S^t with no sure-exits (i.e., $p_s<1$ for all s).

Taking $t_0>>\max\{t:S_0(t)>0\}$, consider a deviation by the proposer from the equilibrium strategy x to $\tilde{x}$ given by $\tilde{x}_t(S_0|t_0)=x_t(S_0|t_0)+\nu$ and $\tilde{x}\equiv x$ elsewhere. In words, the proposer commits to demand ν more at time t_0 if receiver has been playing according to S_0 until then; the original strategy is otherwise unchanged. We will argue that this deviation is profitable if we appropriately choose ν small enough and t large enough, leading to a contradiction.

The proposer's direct gain from this deviation is at least $\delta_1^{t_0} P(S_0)\tilde{P}(S_0)\nu$, where $\tilde{P}(S_0) = \prod_{t=0}^{\infty}(1-p_t) > 0$ is the probability that the receiver stays in the game forever (hence until t_0) conditional on playing S_0. Denoting by $M = \frac{\bar{\theta}+\eta}{1-\delta_1}$ the proposer's maximum continuation value at any point (and hence her maximum loss from the receiver exiting), the loss from this deviation is at most $M \times (\sum_{s\leq t} \delta_1^s Q_s)$, where Q_s is the probability that the receiver changes her choice in period s as a result of the proposer's deviation. Formally, $Q_s = \sum_{S^s} Q_s(S^s)$, where S^s runs over all possible length-s receiver action paths, and $Q_s(S^s)$ is the probability that the receiver changes her action in period s in response to the deviation conditional on having played S^s before period s, i.e.,

$$
\begin{aligned}
Q_s(S^s) = \sum_{i=0}^{k} |&\Pr(\text{receiver plays } p^i \text{ at } s | \text{receiver played } S^s \text{ and faces } \tilde{x}) - \\
&\Pr(\text{receiver plays } p^i \text{ at } s | \text{receiver played } S^s \text{ and faces } x)|.
\end{aligned}
$$

Note that, while in the baseline model the receiver might quit in response to an escalation, here the receiver may respond by quitting outright, or increasing or decreasing her quitting probability, i.e., switching from one p^i to another; we bound the proposer's loss from any such action switch by M. Note also that, of course, $Q_s \equiv 0$ for $s > t_0$, since the receiver's incentives are unchanged conditional on reaching period $t_0 + 1$ with the same action path. And, for $s \leq t_0$, $Q_s(S^s) = 0$ if $S^s \neq S_0|s$: since $\tilde{x}$ and x only differ in the demand made after the receiver plays the action path $S_0|t_0$, the receiver's incentive are the same under $\tilde{x}$ and x if the receiver has already deviated off this path.

Hence we only need to bound $Q_s(S_0|s)$ for $s \leq t_0$. We proceed in similar fashion to the proof of Proposition 4. Briefly, choose $m < t_0$ fixed. For $s = t_0 - m, \dots, t_0$, we use the fact that, if s is large enough, almost no receivers would quit in the absence of a deviation, since by assumption $P(S_0|s) \searrow P(S_0) > 0$ and, by construction, $\tilde{P}(S_0|s) = \tilde{P}(S_0)$ for $s > \max\{t : S_0(t) > 0\}$. As a result, the remaining receivers' continuation value is nearly flat as a function of θ for large s, whereas the continuation payoff from deviating to some $p_s > 0$ (hence $\geq p^1$) varies at a rate equal to at least $\frac{1}{\delta_2}p^1$ in absolute value, as a function of θ. Thus, since even the lowest receiver types left in the support are at least marginally willing to choose $p_s = 0$ on path, it follows that all higher receivers are sure stayers, even if a little more is taken. In similar fashion to Proposition 4, we can show that $\sum_{s=t_0-m}^{t} Q_s(S_0|s)$ converges to an expression of the form $O(\nu^2)$ as $t_0 \to \infty$¡ for any fixed m.

On the other hand, for $s < t_0 - m$, we use that the receiver prefers p^i to p^j iff a condition of the form $W_i - W_j + \epsilon_s(p^i - p^j) > 0$ holds. (This is true simply because the receiver's payoff from choosing $p_s = p^i$ is $(1-p^i)(-x_s - \epsilon_s + \delta_2 V_i) - p^i \frac{\theta}{1-\delta_2}$, which has derivative p^i with respect to ϵ_s.)

If p^i is the action taken in period s by S_0, then W_i is a function of ν (with bounded derivative, bounded by an expression of the form $D\delta_2^{t_0-s}$); otherwise, W_i is independent of ν. Then $Q_s(S_0|s) \leq k\frac{D\delta_2^{t_0-s}}{\min|p^{i+1}-p^i|}\overline{g}\nu$. Picking m large enough, we can make these terms smaller than the gain, using that $\delta_2 < \delta_1$.

In the no commitment cases (ii) and (iii), the result is much simpler: as in the original model with binary receiver action, if the receiver stays in the game forever with positive probability, an increase of ν in the current demand goes unpunished with probability going to $1 - o(\nu)$ as $t \to \infty$, under assumption A2(η). $\square$

References

Acharya, Avidit and Juan Ortner, "Progressive Learning," *Econometrica*, 2017, *85* (6), 1965–1990.

Ali, S. Nageeb, Navin Kartik, and Andreas Kleiner, "Sequential Veto Bargaining With Incomplete Information," *Econometrica.*

Baliga, Sandeep and Tomas Sjöström, "Long Wars," 2023. Working paper.

Board, Simon and Marek Pycia, "Outside Options and the Failure of the Coase Conjecture," *American Economic Review*, 2014, *104* (2), 656–671.

Coy, Peter, "What Game Theory Says About China's Strategy," 2021. `https://www.nytimes.com/2021/09/13/opinion/game-theory-china.html`.

Deneckere, Raymond and Meng-Yu Liang, "Bargaining With Interdependent Values," *Econometrica*, 2006, *74* (5), 1309–1364.

Evdokimov, Kirill S., "The Coase Conjecture and Agreement Rules in Policy Bargaining," 2023. Working paper.

Fearon, James D, "Rationalist Explanations for War," *International Organization*, 1995, *49* (3), 379–414.

Fearon, James D., "Bargaining Over Objects That Influence Future Bargaining Power," 1996. Mimeo.

Feinberg, Yossi and Andrzej Skrzypacz, "Uncertainty About Uncertainty and Delay in Bargaining," *Econometrica*, 2005, *73* (1), 69–91.

Fey, Mark, Adam Meirowitz, and Kristopher W Ramsay, "Credibility and Commitment in Crisis Bargaining," *Political Science Research and Methods*, 2013, *1* (1), 27–52.

_ **and Brenton Kenkel**, "Is an Ultimatum the Last Word on Crisis Bargaining?," *The Journal of Politics*, 2021, *83* (1), 87–102.

_ **and Kristopher W Ramsay**, "Uncertainty and Incentives in Crisis Bargaining: Game-Free Analysis of International Conflict," *American Journal of Political Science*, 2011, *55* (1), 149–169.

Freixas, Xavier, Roger Guesnerie, and Jean Tirole, "Planning Under Incomplete Information and the Ratchet Effect," *The Review of Economic Studies*, 1985, *52* (2), 173–191.

Fuchs, William and Andrzej Skrzypacz, "Bargaining With Arrival of New Traders," *American Economic Review*, 2010, *100* (3), 802–836.

_ and _ , "Bargaining with Deadlines and Private Information," *American Economic Journal: Microeconomics*, 2013, *5* (4), 219–243.

_ and _ , "Bridging the Gap: Bargaining With Interdependent Values," *Journal of Economic Theory*, 2013, *148* (3), 1226–1236.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson, "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 1986, *39* (1), 155–190.

Hart, Oliver D and Jean Tirole, "Contract Renegotiation and Coasian Dynamics," *The Review of Economic Studies*, 1988, *55* (4), 509–540.

Himmelman, Jeff, "A Game of Shark and Minnow," 2013. `https://www.nytimes.com/newsgraphics/2013/10/27/south-china-sea/index.html`.

Hwang, Ilwoo and Fei Li, "Transparency of Outside Options in Bargaining," *Journal of Economic Theory*, 2017, *167*, 116–147.

Krause, Peter and Ehud Eiran, "How Human Boundaries Become State Borders: Radical Flanks and Territorial Control in the Modern Era," *Comparative Politics*, 2018, *50* (4), 479–499.

Laffont, Jean-Jacques and Jean Tirole, "The Dynamics of Incentive Contracts," *Econometrica*, 1988, *56* (5), 1153–1175.

Mearsheimer, John J., "Playing with Fire in Ukraine," *Foreign Affairs*, 2022, *17*. `https://www.foreignaffairs.com/ukraine/playing-fire-ukraine`.

Milgrom, Paul and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 2002, *70* (2), 583–601.

Ortner, Juan, "Durable Goods Monopoly With Stochastic Costs," *Theoretical Economics*, 2017, *12* (2), 817–861.

Powell, Robert, "War as a Commitment Problem," *International Organization*, 2006, *60* (1), 169–203.

Romer, Thomas and Howard Rosenthal, "Political Resource allocation, Controlled Agendas, and the Status Quo," *Public Choice*, 1978, pp. 27–43.

_ **and** _ , "Bureaucrats Versus Voters: On the Political Economy of Resource Allocation by Direct Democracy," *The Quarterly Journal of Economics*, 1979, *93* (4), 563–587.

Saleh, Mohamed, "On the Road to Heaven: Taxation, Conversions, and the Coptic-Muslim Socioeconomic Gap in Medieval Egypt," *The Journal of Economic History*, 2018, *78* (2), 394–434.

_ **and Jean Tirole**, "Taxing Identity: Theory and Evidence From Early Islam," *Econometrica*, 2021, *89* (4), 1881–1919.

Schelling, Thomas C., *Arms and Influence*, Yale University Press, 1966.

Schwarz, Michael and Konstantin Sonin, "A Theory of Brinkmanship, Conflicts, and Commitments," *The Journal of Law, Economics, & Organization*, 2008, *24* (1), 163–183.

Tirole, Jean, "From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection," *Econometrica*, 2016, *84* (4), 1291–1343.