Outside Options and Optimal Bargaining Dynamics

Andrew McClellan*
University of Chicago Booth School of Business

May 1, 2020

Abstract

This paper studies a bargaining game between two players, P and A, where A's outside option is stochastic and changes over time. We use a mechanism design approach to solve for optimal bargaining strategies, and find that a new, but intuitive, set of dynamics arise. When A's outside option increases, A is tempted to cease bargaining. To prevent this response, P increases A's continuation value via two means: directly by promising A a larger share of the pie (decreasing demands) and indirectly by giving A more time to explore his outside option before being forced to make a decision (decreasing pressure). We show this solution can be implemented without commitment using either alternating offers bargaining or a simple type of contract we define, an *option with escape clause*.

1 Introduction

When firms and workers negotiate over wages, the worker's outside options are crucial for determining what wage offers he will accept. As negotiations go on, the outside option to the worker may change as he acquires new outside offers or the market demand for his skill set changes. If the worker delays accepting the firm's offer and his outside option becomes better, previous wages offered by the firm may no longer be enough for the worker to accept. How should the offers the firm makes depend on changes in the

We would like to thank Eric Budish, Jack Fanning, Rohit Lamba, Elliot Lipnowski, Doron Ravid and Jakub Steiner as well as seminar audiences at Rochester and Northwestern for helpful discussions and comments. This work is supported by the Liew Family Junior Faculty Fellowship at the University of Chicago Booth School of Business.

* Email: andrew.mcclellan@chicagobooth.edu.

outside option? Should the firm make a take-it-or-leave-it offer to the worker or give his time to explore his outside option? These questions point toward two fundamental aspects of negotiation: how much to *demand* (i.e., how high or low to make the wage) and how much *pressure* to apply (i.e., how long to let the worker consider the offer).

In this paper, we study negotiation dynamics by analyzing a classic "split-the-pie" bargaining problem between two players, P and A, to which we add a changing outside option for player A.[1] We characterize the P's optimal bargaining strategy find that delay and bargaining breakdowns are both prevalent *and* efficient. P's optimal bargaining process features a history dependence that resembles haggling: when A's outside option is high, A can threaten to quit bargaining, leading P to lower his demands. We also find that the pressure exerted by P decreases as well, giving A more time to explore his outside options before being forced to either agree to a split or take the outside option. Our results show a *complementarity* between demands and pressure in how P uses them to provide incentives to A.

Unlike much of the bargaining literature, our results generate efficient delay and gradual concessions in bargaining demands. To understand the intuition for why delay may be efficient, we can view delay in reaching an agreement as "experimentation" for A and the option value of agreeing to split the pie as "insurance" for the risk of experimentation. Consider a worker deciding whether to take an offer today. If he waits until tomorrow, his outside option may go up or down. If he expects to receive a reasonable offer from the firm tomorrow, he may prefer to delay, knowing he can take the offer tomorrow if his outside option goes down. In this way he can enjoy the benefits when the outside option increases, but still be protected against the risk that it decreases. Thus, even though the firm finds delay purely inefficient, it may benefit from allowing the worker to, at times, explore his outside option if the firm can appropriate a larger part of the surplus by decreasing its offer when the outside option is low.

Although this intuition points out the efficiency of delay, it doesn't tell us what an efficient bargaining process should be. In the first part of the paper, we answer this question by studying a mechanism design problem in which we allow P to commit to his offer process. Because P cannot stop A from taking his outside option, his choice of mechanism must ensure A prefers to continue bargaining at each moment. Treating A's choice to take the outside option early as a deviation, we must consider deviations in an infinite-dimensional space, making analysis of the problem difficult. We identify a binding

[1] Changing outside options is a common feature of many bargaining environments. In firm-union negotiations, the expected payoff to going on strike may change. In buyer-seller negotiations, the value of leaving the negotiation to search for another trading partner is not constant over time; for example, whether it is a "buyers' market" or "sellers' market" may change.

class of constraints on deviations for A and a tractable relaxed problem incorporating only these constraints, which yields a solution to our full problem.

The optimal offer process, although it features non-stationary dynamics, is still simple and intuitive. It can be characterized by a split amount, a split threshold and a breakdown threshold. A split is made when A's outside option goes below the split threshold. The placement of the split threshold tells us how much pressure is placed on A: the higher the threshold, the less time A has to explore his outside options before being forced to accept a split or walk away. We find both the demand by P and the location of the split threshold change over the course of the game, monotonically decreasing in the best outside option that A has received in the past. When A's outside option reaches a new high, P gradually lowers his demand and the pressure on A, keeping it fixed until A's outside option again reaches a new high or an agreement is reached. P's offer process thus features a type of downward rigidity. The bargaining process will not always end with players reaching an agreement: if A's outside option goes above the breakdown threshold, A walks away and takes his outside option. We find that these dynamics persist when considering a social planner designing the optimal mechanism and we map out the payoff frontier of the game. Besides matching natural bargaining dynamics, our results also give a new relationship between players' outside options and split amounts and delay; for example, the higher past or current outside options have been, the longer delay until a split is reached and the agreed upon split will be larger.

In the second part of the paper, we examine whether we can weaken P's commitment power by looking at two related discrete-time bargaining games. In our first exercise, we design a simple type of contract, an *option with escape clause*, that require only minimal commitment power and, in the continuous time limit, implement our mechanism design solution. This contract gives A the option to request a prespecified split at the time of his choosing but allows P to pay an escape penalty to cancel the contract. These contacts are simple and do not require an outside enforcement body to observe A's outside option. When P can use such contracts, we show a *unique* stationary equilibrium exists and that the equilibrium outcome converges to our mechanism design solution as the period length becomes small.

In our second exercise, we look at a classic discrete-time alternating-offers bargaining game in which players can use "short-lived" offers. For every point on the Pareto-frontier, we construct equilibria that converge that point in the frequent offer limit. These equilibria are built by approximating our mechanism-design solution, and thus retain the same dynamics as described above. Together, these two results show the loss from dropping commitment is negligible.

2 Related Literature

Bargaining is an important aspect of many economic interactions and has received considerable attention within the game-theory literature. Rubinstein (1982) established the uniqueness of equilibrium outcomes in an infinite-horizon alternating-offers bargaining model and found that an agreement is reached immediately. The finding of no delay in reaching an agreement is at odds with some real-world phenomena (e.g., haggling, labor strikes, etc.). Several strands of literature have explored reasons for delay in bargaining, such as incomplete information (Fudenberg and Tirole (1985), Gul and Sonnenschein (1988)) or reputational incentives (Abreu and Gul (2000)).

Our paper is broadly related to strands in the bargaining literature looking at the role of players' outside option and the role of changing bargaining environments. The importance of the outside option in bargaining is well known and has been studied in axiomatic bargaining (Nash (1950)), strategic bargaining (Binmore and Sutton (1989)), in conjuction with reputation (Compte and Jehiel (2002), Lee and Liu (2013)) and in relation to the Coase conjecture (Board and Pycia (2014)). These papers assume players' outside options stay fixed throughout the game. The fact that the outside option is dynamic in our model ties us to the literature on bargaining in changing envrionments, which has received growing attention in recent years. This literature has looked the impact of newly arriving players (Fuchs and Skrzypacz (2010), Chaves (2019)), the impact of transparency of outside options (Hwang and Li (2017)), the arrival of information about a seller's types (Daley and Green (2018)) and changing costs of supplying a good (Ortner (2017)). These papers have focused on studying stationary equilibria (where players' strategies depend on beliefs about their opponent's type) and find that incomplete information about players' preferences may generate delay.

Our paper fundamentally differs from much of the bargaining literature in that our main result takes a dynamic contracting approach to the problem by allowing one party to commit to their offer process; only later do we explore equilibrium in our environment. This approach allows us to solve for efficient bargaining outcomes and show they possess a relatively simple structure. Our optimal mechanism generates dynamics with haggling and breakdowns, as are often observed in real-world negotiations. The structure of offers and delay generated in our paper features sporadic concessions and periods of intrasigence. Papers in the Coasian bargaining literaure, such as Fuchs and Skrzypacz (2010), generate a cream-skimming style of delay, finding equilibrium with a gradual, but deterministic, downward movement in offers, whereas papers in the reputational literature, such as Abreu and Gul (2000), generate a war-of-attrition style of delay, finding equilibrium in which any concession in bargaining demand leads to immediate agreement.

Although driven by different forces than our own, efficient delay may arise in Merlo and Wilson (1995) and Cripps (1998), who study models where the size of the surplus to be split is stochastic. In these models, players may benefit from delay only if the expected discounted total surplus tomorrow is greater than the surplus today. Our paper, by contrast, has the outside option changing and assumes the expected discounted value of the outside option is always smaller than the outside option today. Thus, delay in our model is driven by the interplay between the changing outside option and bargaining rather than changes in the outside option alone. Changes in the outside option give a very different set of incentive constraints for players and generate starkly different dynamics than changes in the size of the surplus.

The dynamics in efficient offer processes feature a backloading of incentives as in Ray (2002) and a downward rigidity to P's demands. This type of rigidity is also found in Harris and Holmstrom (1982), who find such rigidity in wages that arise from a competitive market for workers, Thomas and Worrall (1988), who study the design of self-enforcing contracts, and in McClellan (2019), who studies the design of approval rules to incentivize experimentation.

3 Model

Two players, P and A, bargain over how to split a pie of size one. P and A have utility functions u_P and u_A over the share of the pie they receive when an agreement is reached, at which point the game ends. Time runs continuously from $t = 0$ to ∞, and both players discount time at a rate of $r > 0$.

Each player has an outside option that they can take at any time, immediately ending the game. Both players will recieve their outside options in this event. The dynamics in our model are driven by A's outside option $X_t \in [\underline{X}, \overline{X}]$. Starting at X_0, the evolution of X_t is given by the stochastic differential equation

$$dX_t = \mu(X_t)dt + \sigma(X_t)dB_t,$$

where $B = \{B_t, \mathcal{F}_t, 0 \leq t \leq \infty\}$ is a standard Brownian motion on the canonical probability space $(\Omega, \mathcal{F}, \mathcal{Q})$ subject to standard conditions. We assume both $\mu(\cdot), \sigma(\cdot)$ are Lipschitz continuous and $\sigma(\cdot) > 0$ on $(\underline{X}, \overline{X})$. The entire path of X_t is public and observed by both players. A history h_t will consist of the path of X_s from 0 to t. To simplify our main model, we set P's outside option to 0.[2]

To describe the bargaining process, we can heuristically think of P as making an offer at each instant of time, which A can either accept or reject. Our main result focuses on the case in which P is allowed to commit to his offer strategy. We call P's choice of an offer strategy a mechanism, which we define by the outcome it induces.

Definition 1. *A **mechanism** consists of $\mathcal{F}_t$-measurable functions $(\tau, d_\tau, \alpha_\tau)$ where*

1. *τ is a stopping time that gives the time when the game ends; that is, a split is made or either player takes his outside option.*

2. *$d_\tau \in \{0, 1\}$ is a decision rule that equals 1 if and only if a split is made at time τ.*

3. *$\alpha_\tau \in [0, 1]$ gives P's share if a split is made at time τ.*

The expected payoff to P from the mechanism $(\tau, \alpha_\tau, d_\tau)$ is

$$J(\tau, d_\tau, \alpha_\tau, X_0) := \mathbb{E}[e^{-r\tau} d_\tau u_P(\alpha_\tau)|X_0],$$

whereas the expected payoff to A from the mechanism is

$$V(\tau, d_\tau, \alpha_\tau, X_0) := \mathbb{E}[e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0].$$

Without loss, we focus on mechanisms in which A never takes the outside option.[3] Because P can not stop A from taking his outside option, P will need to ensure that A's continuation value in the mechanism remains above his outside option. We formally introduce this constraint this in Section 5.

We place several relatively weak assumptions on the primitives of the model. Our first assumption imposes conditions on the utility functions.

Assumption 1. *The utility functions u_A, u_P are twice differentiable with $u_i'' \leq 0 < u_i'$ for $i = P, A$ (with a strict concavity for some i), $u_A(1) \in (\underline{X}, \overline{X})$ and $0 = u_P(0) = u_A(0) < \underline{X}$.*

The concavity assumption simplifies the derivation of the optimal mechanism and ensures players can't benefit from randomization over splits, whereas the other assumptions allow us to rule out uninteresting cases; for example, if $u_A(1) < \underline{X}$, a split will never be agreed to and A will always take his outside option.

[2] We extend the model to a richer set of outside options for P in Section 7.

[3] Take any mechanism which, after some history, calls for A takes the outside option. We can replace A taking the outside option with P. This will still give the same payoffs to both players.

Our next assumption concerns the evolution X_t. We assume the expected future discounted value of the X_t is lower than the current X_t. This assumption will allow us to ensure that, if A knows there is no possibility of reaching an agreement in the future, then A has no incentive to continue bargaining and would be better off taking the outside option immediately.

Assumption 2. *$e^{-rt}X_t$ is a strict supermartingale.*

Assumption 2 is a natural property to impose on the outside option. Consider a firm-worker wage negotiation where X_t represents the value of searching for new job offers. Because A always has the option to ignore incoming job offers, he can do no worse by reentering the search market immediately rather than continuing bargaining only to reenter the search market in the future with probability one.

3.1 Discussion

Two features in our model worth discussing are the observability of the outside option and the ability of P to commit to his offers. The assumption of common knowledge of X_t is similar to other papers in the literature on changing bargaining environments and is economically reasonable in many situations. If we think of a worker and firm bargaining in tight-knit industries, the offers that a worker has at other firms often can be verifably disclosed by the worker. If a labor union and company are bargaining, X_t may be a measure of how favorably public opinion would view the union if they were to go on strike, which would affect how likely politicians are to intervene in favor of or against the strike.[4] The common-knowledge assumption is also needed for the more pragmatic reason that without it, the model is intractable: without observability of X_t, P must rely on A's report about X_t and the optimal mechanism will need to keep track of P's belief about X_t if A were to have misreported X_t, which will be a complicated object.

In contrast to much of the bargaining literature, we allow P to commit to his offer process. In the case of a firm bargaining with a worker, the firm may have reputational concerns that allow it to commit to its negotiating stance: the firm will be bargaining with many workers over time, and sticking to its bargaining demands today may affect bargaining outcomes in the future. Solving the problem with commitment will prove useful when we study equilibria in discrete-time versions of our model. The commitment solution gives us an upper bound on P's payoffs in any equilibrium of a bargaining game,

[4] In this interpretation, the assumption that P's outside option also depends on X_t seems natural. We study this in Section 7, in which we provide assumptions under which the fundamental structure of our results will not change.

which would otherwise be difficult to solve for. Having identified the upper bound and the offer process which achieves it, it will be much easier to construct equlibria which reach this upper bound as we approach the continuous time limit.

4 Benchmarks

As in most of the bargaining literature, delay in agreeing to a split is purely inefficient. A split that is enacted in the future would be better for both players if it were enacted immediately. By Assumption 2, we know that delay in taking the outside option is also inefficient. It seems natural to conjecture that a Pareto-efficient outcome features either an immediate split or immediately taking the outside option. The economic intuition for why this conjecture is wrong can be seen by viewing delay as A experimenting with his outside option and the option to accept a split of the pie as insurance against a decrease in the outside option. The option of making a split with P adds option value for A from continuing to bargain.

Consider the case when $dX_t = dB_t$.[5] If $X_0 = u_A(1)$, the only possible bargaining split that achieves no delay and respects A's individual rationality is to give the entire pie to A. Is this no-delay outcome efficient?

Consider the alternative offer by P in which he asks A to wait for Δ length of time and commits to give $1 - \Delta^2$ of the pie to A. This new offer features both delay and a lower value of the split for A. If A waits and his outside option goes up, he can take his new higher outside option, but if his outside option goes down, he can take the split if $X_\Delta < u_A(1 - \Delta^2)$. This option value protects him against a decrease in his outside option. A's expected utility of waiting is equal to

$$(1 - r\Delta)u_A(1) + \frac{\sqrt{\Delta}}{\sqrt{2\pi}} + o(\Delta^{\frac{3}{2}}).$$

For small Δ, this policy yields a higher value than the value of stopping immediately, $u_A(1)$, implying ending the game immediately cannot be efficient. The driving force for this result is the fact that allowing A to choose the max over $u_A(1), X_t$ creates a kink in the underlying payoff. The convexity this kink creates is enough, at least when starting close to $X_t = u_A(1)$, to make it beneficial for A to delay and take a lottery over payoffs tomorrow.

[5]Formally, to satisfy Assumption 1, we should require that X_t has some reflecting or absorbing barrier at $\underline{X} > 0$. However, this requirement is not necessary for our example.

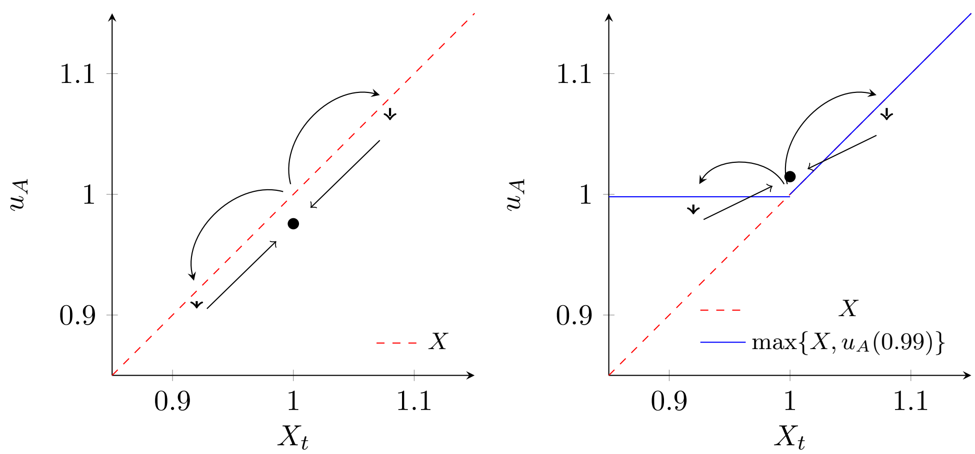

Figure 1: For illustrative purposes, we treat the movement of X_t as a random walk. The upper curved arrows indicate the movement of X_0 to X_Δ, and the downward arrows indicate discounting costs to $e^{-r\Delta}X_\Delta$. The black dot indicates the expected value of waiting until $t = \Delta$ to make a decision. In the left panel, we see this is strictly below the value of the outside option today, whereas in the right panel, using the split as insurance against a decrease in X, we see the value of waiting is higher than the outside option today.

This intuition is economically relevant in many bargaining situations. Coming back to our example of a firm negotiating a wage with a worker, the firm might be able to make a take-it-or-leave-it offer that the worker would choose to accept. However, this offer may require such a high wage that the firm may prefer to give the worker a lower offer but grant the worker time to explore his other options before deciding whether to accept the firm's offer. Such delay is natural in labor markets, where workers are often given offers that do not immediately explode, allowing the worker to take time to explore employment options at other firms.

The argument above shows the benchmark of no delay is not always an efficient outcome. Another natural benchmark we might consider is that of a social planner who places $\zeta, 1-\zeta$ weight on P, A's utility, respectively. If we ignore the possibility that A will take the outside option early, the social planner will then choose a mechanism that solves

$$\sup_{(\tau,\alpha_\tau,d_\tau)} \mathbb{E}[e^{-r\tau}\bigg(d_\tau\big(\zeta u_P(\alpha_\tau) + (1-\zeta)[u_A(1-\alpha_\tau) - X_\tau]\big) + (1-\zeta)X_\tau\bigg)|X_0]$$

The solution takes the familiar form of a *stationary policy*.

Proposition 1. *There are* $(b, B, \alpha_s) \in \mathbb{R}^3$ *with* $b < B$ *such that the optimal policy takes the form*

$$\tau = inf\{t: \ X_t \notin (b, B)\}, \ d_\tau = \mathbb{1}(X_\tau = b), \ \alpha_\tau = \alpha_s.$$

This stationary structure is familiar from standard solutions in single decision maker stopping problems.[6] However, the social planner's problem doesn't take into account the incentive constraint that A must find it optimal to delay taking the outside option until the prescribed time. We argue that any stationary mechanism in which P gets positive utility and respects A's incentive constraints *cannot* be efficient.

Consider a stationary mechanism that respects A's incentive constraints and calls for P to offer α_s. Let's go to the moment when X_t has reached B and A is about to take his outside option. If A does so, P will go home with a payoff of 0. Imagine if P were to come to A and propose a new continuation mechanism in which, for some small ϵ, P always demands $\alpha_s - \epsilon$ and lets A choose when to accept the split or take his outside option. This would increase the value of bargaining at X_t, and so if A were indifferent between continuing and taking his outside option at X_t, he would strictly prefer to continue bargaining. Moreover, this would also increases P's utility because A might eventually agree to a split. This argument implies any stationary policy for which $\alpha_s > 0$ can be improved upon. Understanding how to best design bargaining strategies when faced with A's incentive constraints is the subject of the next section.

5 Mechanism Design Problem

We now turn to the problem of how to design the optimal mechanism. The key constraint P faces is a dynamic interim-rationality constraint that ensure that A doesn't have an incentive to take his outside option early. Suppose that A were to follow a strategy in which he took his outside option early at some τ'. His expected payoff from continuing to taking the outside option at τ' would be $V(\tau \wedge \tau', d_\tau \mathbb{1}(\tau < \tau'), \alpha_\tau, X_0)$. Our mechanism will need to ensure that, for every τ', A could choose, his value of following the mechanism is weakly higher. Formally, we write our constraint on P's choice of mechanism, which we will call DIR, as[7]

$$(DIR): \sup_{\tau' \in \mathcal{T}} V(\tau \wedge \tau', d_\tau \mathbb{1}(\tau < \tau'), \alpha_\tau, X_0) \leq V(\tau, d_\tau, \alpha_\tau, X_0),$$

[6] The proof is a straightforward application of dynamic programming and is hence omitted.

[7] DIR is slightly weaker than the stronger interim rationality constraint that A's continuation value is greater than the outside option after every history. In Lemma E.13, we show that any mechanism for which A's continuation value is greater than the outside option satisfies DIR. After solving the problem with only DIR, we will verify that the solution satisfies the stronger interim rationality constraint.

where $\mathcal{T}$ is the set of all $\mathcal{F}_t$-measurable stopping rules. TP can ensure A does not delay in accepting an individually rational offer at time τ by committing to demand $\alpha_t = 1$ thereafter in the event that A does not agree to the split at time τ.

We also add a promise-keeping constraint, PK, that ensures that A receives at least W expected utility from the mechanism (in the case of P's optimal mechanism, we can take $W = X_0$). This constraint does not add any additional difficulty to our mechanism-design problem and proves useful later when characterizing the payoff frontier:

$$PK(W): \; V(\tau, d_\tau, \alpha_\tau, X_0) \geq W.$$

Adding this promise-keeping constraint also allows us to consider situations in which a hold-up problem exists: if A must make some costly investment prior to the start of bargaining, the fact that P's mechanism ensures A has enough continuation value to find making the costly investment profitable is important. This type of situation arises naturally in our motivating examples.[8]

Let $V^*(X_0) := \sup_{\tau, d_\tau} V(\tau, d_\tau, 1, X_0)$, which is the highest utility A could receive in any mechanism (we call its solution the A-*optimal mechanism*). To ensure a solution that satisfies $PK(W)$ exists, we assume $W \in [X_0, V^*(X_0)]$. We can then formally state our mechanism problem:

$$J^*(X_0; W) := \sup_{(\tau, d_\tau, \alpha_\tau)} J(\tau, d_\tau, \alpha_\tau, X_0) \qquad (1)$$
$$\text{subject to } DIR, \; PK(W).$$

Even though we allow for arbitrarily complex mechanisms, the optimum still turns out to be quite simple and intuitive. The optimal mechanism is measurable with respect to only two state variables, X_t and the running maximum $M_t := \max_{s \in [0,t]} X_s$, and can be described by three objects: an offer function $\alpha(M_t)$, a split threshold $S(M_t)$, and a breakdown threshold $\overline{R}$. An agreement to split the pie is reached whenever $X_t \leq S(M_t)$ and A receives his outside option if and only if $\overline{R}$ is reached before $S(M_t)$. We interpret the location of the split threshold S as the amount of *pressure* being placed on A: the higher S is, the less time A has to explore his outside option before being forced to make a decision (i.e., pressure is higher). By looking at how α, S change with M_t, we can see how the demands and pressure change over the course of the game. Theorem 1 gives properties of the optimal mechanism.

[8] In the case of a firm-worker negotiation, the worker may have to expend effort to develop firm-specific human capital that creates the surplus over which the players are bargaining. If we think about a buyer-seller negotiation, the buyer may have to expend effort to learn the value of the seller's good (i.e., the size of the pie) or incur search costs to discover the seller.

Theorem 1. *There optimal mechanism* $(\tau^*, \alpha^*_\tau, d^*_\tau)$ *is given by, for some decreasing continuous functions* $S(\cdot), \alpha(\cdot)$*,*

$$\tau^* = inf\{t : \; X_t \notin (S(M_t), \overline{R})\},$$
$$d^*_\tau = \mathbb{1}(X_\tau \leq S(M_\tau)),$$
$$\alpha^*_\tau = \alpha(M_\tau).$$

The proof of Theorem 1, along with all other proofs, is in the Appendix. Before providing a sketch of the proof, we first discuss several features of the optimal mechanism. Intuitively, when X_t reaches a high, A is tempted to take the outside option and P must increase A's continuation value to prevent A from walking away. We show that P optimally increases A's continuation value by decreasing his own demand (thereby increasing A's utility in the event that they agree to a split) *and* decreasing the pressure on A (i.e., lowering the threshold at which A takes the split). Both of these changes are rigid, never rising once they have decreased, implying a persistent effect from A having higher outside options in the past.

The optimal mechanism features a very simple history dependence, only relying on M_t. Although P *could* use more complex schemes to increase utility, our results show a simple mechanism does better. If P were to increase his demand in the future, he would have to decrease it even more today in order to deliver the necessary continuation value to A today. P instead finds it optimal to smooth the decrease in his demand over time as long as he can, adjusting demand only when A is tempted to walk away. This smoothing motive gives the decrease in α, S a downward rigidity.

Once we know history dependence will exist in the optimal mechanism, the decreasing demand α is natural. If A has a better outside option, P has to offer A a larger split of the pie in order prevent A from taking his outside option. This intuition, although correct, turns out to be a bit more nuanced once we note that decreasing α is not P's only way to increase A's continuation value. Because A can benefit from exploring his outside option, lowering S provides another lever by which P can increase A's continuation value.

Increasing A's utility via a lower S, although costly for P in terms of discounting, allows P to lessen the decrease in α that would otherwise be necessary if, say, the threshold S were stationary. Our result shows the demands and pressure from P are *complements*: the higher the demand α, the higher P's marginal utility from raising S. Thus, to increase A's continuation value, P will optimally use both tools together, jointly decreasing demands and pressure.

Additionally, we find that bargaining breakdowns, when A chooses to take his outside option, will always happen with positive probability when delay in the optimal mechanism exists. This result stands in contrast to the stationary bargaining literature, where players will always agree to a deal on-path. A notable exception is Board and Pycia (2014): they find that in a buyer-seller model some buyer types will choose to take their outside options. Experimentation only yields benefits if the results of the experimentation, a higher X_t, are used. If bargaining never broke down, delay would be inefficient and P would be better off making an offer that A immediately accepts.

The fact that α, S are jointly decreasing has a number of interesting applications for observable outcomes. For example, fixing X_t, the higher A's outside options have been in the past, the longer the delay will be until an agreement is reached. We also see that when an agreement is reached, A receives a bigger split of the pie when his outside option at the time of the split is lower. For an outside observer, these dynamics might appear to be something like anchoring effects, stubbornness or loss aversion. Our results show how such dynamics arise using standard preferences and are necessary for bargaining outcomes on the Pareto-efficient frontier.

Although the mechanism only requires that P makes the offer of $\alpha(M_\tau)$ at time τ, it is without loss that P makes the offer of $\alpha(M_t)$ at each instant of time. Because the $\alpha(\cdot)$ is decreasing over the course of the game, P would be better off if A were to immediately take the offer of α. If taking the split were in A's interest, we would have a Pareto improvement that doesn't violate the DIR constraints, a contradiction. Even if P offers $\alpha(M_t)$ at each instant, A will never find it optimal to take this offer early.

5.1 Proof Outline

We start our sketch of the proof by solving for the optimal form of d_τ. In the Appendix we show that A-optimal mechanism, which solves $V^*(X_0)$, involves him taking the outside option whenever X_t goes above some $\overline{R}$. Whenever $X_t \geq \overline{R}$, any mechanism which satisfies DIR must involve immediately taking the outside option. In the other direction, whenever $X < \overline{R}$, taking the outside option cannot be optimal: if P were to offer to let A take the entire pie at any time of A's choosing, A would find it optimal to delay taking his outside option. This option yields the same utility to P, because P's utility is the same from giving away the pie or letting A taking his outside option. However, P does not need to use such an extreme offer: he could instead offer to let A take $1-\epsilon$ of the pie at any time of A's choosing. For small ϵ, it would still be optimal for A to delay taking the outside option. This modification would improve on the previous mechanism for P because P receives $u_P(\epsilon) > 0$ with positive probability. Therefore, we can restrict attention to mechanisms such that $d_\tau = 0$ if and only if $X_\tau \geq \overline{R}$.

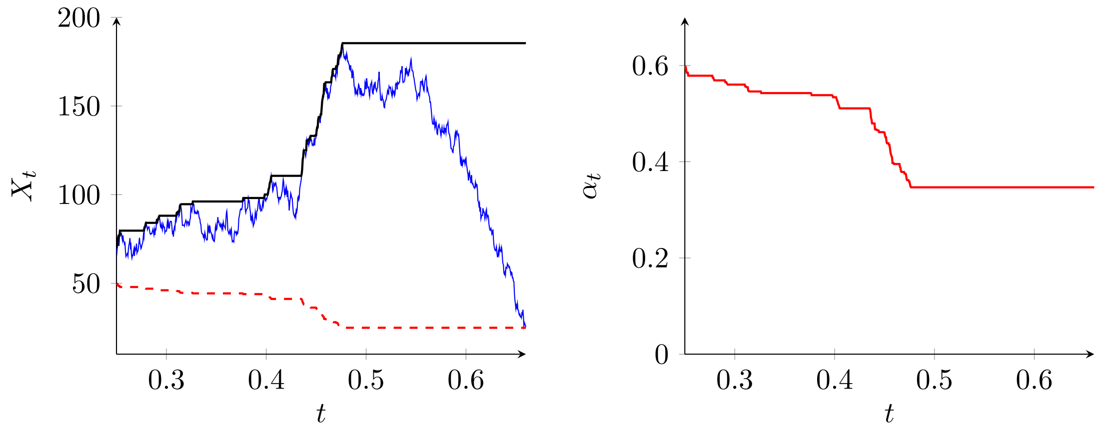

Figure 2: In the left graph, we plot a sample path of X_t in blue, M_t in black and $S(M_t)$ in the red dashed line. In the right graph, we plot the corresponding path of P's demands. We can see the demand and pressure exerted by P are monotonic, decreasing rapidly in spurts and then remaining fixed as X_t goes down.

Lemma 1. *The optimal mechanism that satisfies DIR has* $\tau \leq inf\{t: \ X_t \geq \overline{R}\}$ *and* $d_\tau = 0$ *if and only if* $X_\tau \geq \overline{R}$.

Lemma 1 allows us to focus only on optimizing over (τ, α_τ). For those familiar with the dynamic contracting literature, the most natural route to take would be to treat the agent's continuation value as a state variable and use a dynamic programming approach to solve the problem (e.g., Sannikov (2008)). However, this approach runs into difficulty once we note we also need to keep track of X_t as a state variable. This method then involves solving a PDE, which is not feasible except in very special cases.

The main difficulty we face is that for an arbitrary (τ, α_τ) that P might choose, finding A's best response among τ' in DIR is not feasible. To get around this difficulty, we use a similar approach to that in McClellan (2019), defining a relaxed problem in which we consider a weaker version of the DIR constraint. We limit the set of τ' deviations that A can choose to stopping rules in which he takes the outside option at the first time X_t goes above some threshold R. We call these deviations *threshold-quitting rules*. The threshold-quitting rule associated with a threshold R is

$$\tau_+(R) := inf\{t: \ X_t \geq R\}.$$

We define $d_\tau(R) = d_\tau \mathbb{1}(\tau \leq \tau_+(R))$ to be the decision rule induced by A when he takes the outside option early at $\tau_+(R)$. A relaxed DIR constraint associated with some $R \in \mathbb{R}$

will be

$$RDIR(R) : V(\tau \wedge \tau_+(R), d_\tau(R), \alpha_\tau, X_0) \leq V(\tau, d_\tau, \alpha_\tau, X_0)$$

Let $\mathcal{X}_N = \{X_1, ..., X_N\}$ be a finite collection of thresholds.[9] For an arbitrary $\mathcal{X}_N$, we study the following *relaxed* mechanism-design problem:

$$\sup_{(\tau,\alpha_\tau)} \; J(\tau, d_\tau, \alpha_\tau, X_0)$$
$$\text{subject to } RDIR(X_n) \; \forall X_n \in \mathcal{X}_N, \; PK(W).$$

This relaxed problem allows us to use a Lagrangian approach to convert this constrained optimization problem into a single-decision maker optimal stopping problem.[10] With a set of Lagrange multipliers $(\lambda(X_1), ..., \lambda(X_N)) \in \mathbb{R}^N_-$ associated with the $RDIR(X_n)$ constraints and $\eta \in \mathbb{R}_-$ associated with the PK constraint, the solution to our relaxed problem solves

$$\begin{aligned}\sup_{(\tau,\alpha_\tau)} \; \mathbb{E}\Big[& e^{-r\tau}\big(d_\tau\{u_P(\alpha_\tau) - \eta\big((u_A(1-\alpha_\tau) - X_\tau\big)\} - \eta X_\tau\big) \\ & + \sum_{n=1}^{N} \lambda(X_n)\{e^{-r(\tau\wedge\tau_+(X_n)}\big(d_\tau(X_n)(u_A(1-\alpha_\tau) - X_{\tau\wedge\tau_+(X_n)})\big) + X_{\tau\wedge\tau_+(X_n)})\big) \\ & - e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)\}|X_0\Big].\end{aligned} \tag{2}$$

Let $\mathcal{B} = \{X^1, ..., X^M\}$ be the set of thresholds $X_n \in \mathcal{X}_N$ such that $RDIR(X_n)$ is binding in the solution, ordered $X^1 < X^2 < ... < X^M$; complementary slackness implies $\lambda(X_n) < 0$ only if $X_n \in \mathcal{B}$. Although the Lagrangian in 2 may appear complicated, the conversion to a unconstrained single-decision-maker problem allows us to apply optimal stopping arguments to pin down much of the structure of the solution.

The solution to 2 possesses a kind of "local stationarity" in the stopping rule and split amount. Let us focus on the optimal rule before $\tau_+(X^1)$. Conditional on stopping, we choose α_τ optimally. Prior to $\tau_+(X^1)$, the choice of α_τ solves

$$\begin{aligned}& argmax_{\alpha\in[0,1]} \; u_P(\alpha) - \eta u_A(1-\alpha) + \sum_{n=1}^{N} \lambda(X_n)[u_A(1-\alpha_\tau) + d_\tau(X^n)u_A(1-\alpha_\tau)] \\ & = argmax_{\alpha\in[0,1]} \; u_P(\alpha) - \eta u_A(1-\alpha).\end{aligned}$$

[9] For technical reasons, we restrict attention to a finite grid of threshold quitting rules and then look at the limit as this grid gets arbitrarily fine.

[10] Strong duality and complementary slackness hold by Dokuchaev (1997) and Balzer and Janßen (2002).

Thus, when stopping prior to $\tau_+(X^1)$, the optimal split amount will always be the same. We also note for any $s, t < \tau_+(X^1)$, the continuation value at $X_t = X$ is equal to the continuation value at $X_s = X$ prior to $\tau_+(X^1)$. This observation allows us to argue that, as long as $\tau_+(X^1)$ has not been reached, the decision of whether to stop depends only on the current X. We show there exists a "split threshold" S_0 exists such that stopping is optimal if and only if $X_t \leq S_0$. Because both the threshold and split amount are stationary only as long as we have not reached $\tau_+(X^1)$, we call (τ, α_τ) locally stationary.

After $\tau_+(X^1)$, the structure of τ and α_τ will change. The continuation value at $\tau_+(X^1)$ is equal to

$$\sup_{(\tau,\alpha_\tau)} \mathbb{E}\Bigg[e^{-r\tau}\big(d_\tau[u_P(\alpha_\tau) - (\eta + \lambda(X^1))(u_A(1-\alpha_\tau) - X_\tau)] - (\eta + \lambda(X^1))X_\tau\big) \\ + \sum_{n=2}^{N} \lambda(X_n)\big\{e^{-r(\tau \wedge \tau_+(X_n))}\Big(d_\tau(X_n)\big(u_A(1-\alpha_\tau) - X_\tau\big) + X_{\tau\wedge\tau_+(X_n))}\Big)] \\ - e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)\big\}|X^1\Bigg] + \lambda(X^1)X^1. \tag{3}$$

Note that 3 is independent of the history of play prior to $\tau_+(X^1)$. This independence allows us to show the optimal continuation mechanism at $\tau_+(X^1)$ is the same regardless of the history of play prior to $\tau_+(X^1)$. We apply the same arguments to conclude the optimal rule will again have a locally stationary split amount and threshold (α_1, S_1) until $\tau_+(X^2)$. Applying these arguments repeatedly, the solution to our relaxed problem is given a split amount and thresholds (α_m, S_m) for each threshold X^m. Our next proposition formally states this result along with how they change with respect to m, the intuition for which we discuss next.

Proposition 2. *The optimal* τ, α_τ *that solves* 2 *can be written as follows: for some* $(S_0, ..., S_M) \in \mathbb{R}^{M+1}$ *and* $(\alpha_0, ..., \alpha_M) \in [0,1]^{M+1}$*:*

$$\tau = \tau_+(\overline{R}) \wedge inf\{t : \; M_t \in (X^m, X^{m+1}] \; and \; X_t \leq S_m\}$$

$$\alpha_\tau = \sum_{m=0}^{M} \alpha_m \mathbb{1}(X_\tau \leq S_m).$$

Moreover, α_m *and* S_m *are decreasing in* m*.*

Proposition 2 tells us the optimal mechanism changes whenever a new X^m threshold is reached. Because M_t is a sufficient statistic for which X^m have been reached, we can write the solution using only two state variables, X_t and M_t. We can see α will be decreasing by observing that the more X^m that have been reached, the larger the Lagrange multiplier on A's utility becomes. Compare problems 2 and 3. In the first, P chooses α to maximize $u_P(\alpha)-\eta u_A(1-\alpha)$ while. In the second, he chooses α to maximize $u_P(\alpha)-(\eta+\lambda(X^1))u_A(1-\alpha)$, which has a greater weight on u_A. This argument extends for $X^2, ..., X^M$. The more nuanced argument comes when looking at how S_m changes in response to M_t.

P's payoff is reduced when the split threshold S_m is lower for two reasons: lowering S_m both lengthens the time until a split is reached (generating higher discounting costs) and increases the probability that X_t reaches a new X^m, after which P must lower his demand to provide additional continuation value to A. This observation implies P's optimal choice of S_m will always be higher than the threshold at which A would choose to take the split. If S were strictly below A's preferred threshold, then P could increase S_m and make both players strictly better off. Moreover, the choice of S_m must be *strictly* above A's optimal choice. At A's optimal choice of a split threshold, raising the threshold leads to only a second-order loss for A and can be compensated with a second-order decrease in α. However, for P, the benefit of raising the split threshold is first order, making such a trade-off beneficial for P.

When α is lower, delay is less costly for P: discounting costs decrease when P's utility from a split is lower. If P were to consider raising S, he would need to compensate A by decreasing α. The compensating decrease in α becomes more costly the lower α is for two reasons. First, due to the concavity of u_P, a decrease in α is more costly for P when starting at a lower α. Second, due to the concavity of u_A, a larger decrease in α is needed to increase A's utility when α is low. Therefore, the benefit for P of raising S decreases when starting at a lower α. Whenever he needs to increase A's continuation value, P will find it profitable to use decreasing pressure as way to mitigate the decrease in α that would otherwise be necessary.

To illustrate this intuition formally, consider the choice of the threshold S at some $X < X^m$ both before and after X^m has been reached. Because the optimal mechanism is constant between $\tau_+(X^m)$ and $\tau_+(X^{m+1})$, we know that the continuation value for P will be the same at X^m for all $t \in \left[\tau_+(X^m), \tau_+(X^{m+1})\right)$. Let us call this continuation value $H(X^m)$. At $\tau_+(X^m)$, we show that$A$'s continuation value is equal to X^m.

When X^m is reached, P needs to increase A's continuation value. We show that for $X < X^m$, A's continuation value at X will be higher after $\tau_+(X^m)$ than before $\tau_+(X^m)$. We look at how P's optimal choice of S changes when we increase the utility W he must provide A and fix the continuation value at X^m to be $H(X^m)$ for P and X^m for A.

The utility of both players will depend on the expected discounted probabilities of reaching each threshold. Let $\tau_-(S) := inf\{t : X_t \leq S\}$. For a pair of thresholds $S < X^m$, we define the discounted probability that X^m (S) is reached first when starting at X be Ψ (ψ):

$$\Psi(X^m, S, X) = \mathbb{E}[e^{-r\tau_+(X^M)}\mathbb{1}(\tau_+(X^m) < \tau_-(S))|X],$$
$$\psi(X^m, S, X) = \mathbb{E}[e^{-r\tau_-(S)}\mathbb{1}(\tau_-(S) < \tau_+(X^m))|X].$$

Given W and a threshold S, we can pin down the demand $\alpha(W, S)$ needed to deliver the utility W to A. P's choice of S will maximize

$$\Psi(X^m, S, X)H(X^m) + \psi(X^m, S, X)u_P(\alpha(W, S)).$$

Using the notation $\Psi_S := \frac{\partial \Psi(X^m,S,X)}{\partial S}$ (and similarly for ψ, α), we can write the first-order condition for S as

$$\Psi_S H(X^m) + \psi_S u_P(\alpha) + \psi u_P'(\alpha)\alpha_S = 0. \tag{4}$$

The first two-terms, $\Psi_S H(X^m) + \psi_S u_P(\alpha)$, give the direct benefit of raising S (which is positive by our earlier observation), whereas the third term $\psi u_P'(\alpha)\alpha_S$ gives the compensating change in α needed to ensure the delivery of W expected utility to A.

When M_t reaches a new X^m, P needs to increase A's continuation value (i.e., a higher W). To understand how increasing W changes the marginal returns to raising S, we take the derivative of 4 with respect to W, giving

$$\underbrace{\psi_S u_P'(\alpha)\alpha_W}_{(a)<0} + \underbrace{\psi u_P''(\alpha)\alpha_W \alpha_S}_{(b)\leq 0} + \underbrace{\psi u_P'(\alpha)\alpha_{WS}}_{(c)?}. \tag{5}$$

Because α is smaller, the payoff to taking the split is lower and the discounting costs of waiting until S go down for P. This force, conveyed in the first term (a), decreases the benefit of raising S. To interpret the second term (b), we note that due to the concavity of u_P, P's utility loss from decreasing α is higher when α is low. As we increase W, we must decrease α, thereby making the necessary decrease in α in response to an increase in S more costly for P. The third term (c) depends on the change in responsiveness of α_S to changes in W. Whereas the first and second terms are negative (using the fact that both $\alpha_W, \alpha_S < 0$), the sign on the third term is unclear. We need to understand α_{WS} to be able to evaluate the sign of 7.

Using the functional form of $\alpha(W, S)$, we find α_{WS} can be decomposed into two parts:

$$\alpha_{WS} = \underbrace{\alpha_S \frac{-u_A''(1-\alpha)}{u_A'(1-\alpha)}}_{(d)\leq 0} + \underbrace{\alpha_W \frac{-\psi_S}{\psi}}_{(e)>0}$$

When α is lower, concavity in A's utility means the marginal return on $1-\alpha$ is smaller. To compensate A for the utility loss of raising S, a larger decrease in α is needed. This force appears as term (d) and, as with (b), makes the cost of raising S higher. On the other hand, when α is lower, A's value of experimentation decreases because he is happier to take the split. Decreasing the value of experimentation to A reduces the increase in $1-\alpha$ needed to compensate A for an increase in S, reducing the cost of raising S. Although this term goes in the other direction of $(a), (b), (d)$, we can show it cancels out with term (a), leading the conclusion that equation 7 is negative. In the Appendix, we use the fact that equation 7 is negative to show S_m is decreasing in m, completing the proof of Proposition 2.

Because the relaxed problem delivers an upper bound on the full problem, we have a candidate solution to our full problem by looking at the limit of our relaxed problem as our grid of points in $\mathcal{X}_N$ gets finer and finer. In the proof of Theorem 1 in the Appendix we verify this limit is well defined and that, in limit mechanism, A's continuation value is weakly greater than this outside option after all histories, making it a solution to our mechanism-design problem.

5.2 Renegotiation and Pareto Efficiency

Theorem 1 applies for any W that is feasible. By varying the promised utility W, we can map out the entire frontier of efficient bargaining outcomes. Not only does P's optimal solution possess the dynamics we find, but *every* efficient outcome (subject to DIR) can be generated by similar dynamics.

However, the proof of Theorem 1 doesn't tell us whether the optimal mechanism remains efficient over time. If continuation play at some histories were extremely inefficient, we might be worried that players would have an incentive to renegotiate the mechanism. Proposition 3 shows our optimal mechanism is resistant to such concerns: we find an optimal mechanism with no room for Pareto-improvements after any history, both on- *and* off-path. This optimal mechanism can deter deviations even while retaining Pareto efficiency.

Part of the optimal mechanism relies on punishing A if A ever rejects P's offer when the mechanism calls for A to accept. The harshest punishment P can deliver to A is to reduce his continuation value to X_t. However, the optimal mechanism could do this in a number of ways. For example, P could demand $\alpha_t = 1$ for all future t (inducing A to take the outside option immediately) or make a take-it-or-leave-it offer of $\alpha_t = 1 - u_A^{-1}(X_t)$, inducing A to take the split and get utility equal to his outside option. These punishments, because they are off-path, do not affect the value of the mechanism for P. However, some punishments might be Pareto inefficient (e.g., demanding $\alpha_t = 1$ can be improved upon by demanding $\alpha_t = 1 - u_A^{-1}(X_t)$). Fortunately, we can maintain our harshest punishment and preserve Pareto-efficiency by stipulating that if A deviates, P offers the optimal mechanism that delivers A a continuation value equal to X_t. With this formulation of the optimal mechanism, we can state our efficiency result.

Proposition 3. *There exists an optimal mechanism with on-path strategies as in Theorem 1 which is Pareto-efficient after all histories, both on- and off-path.*

We can also show that in P's optimal mechanism, the choice of α, S is *independent* of the initial X_0 and depends only on M_t. In the proof of Theorem 1, we show that when $X_t = M_t$, A is indifferent between taking the outside option and continuing to bargain. Therefore, whenever $X_t = M_t$, the optimal continuation mechanism will be the same as the optimal mechanism from starting at $X_0 = X_t$. Therefore, the form of the optimal mechanism at a particular (X_t, M_t) is independent of the starting X_0. This independence is a standard feature in individual decision-maker problems, but is not always true when we include strategic interactions between players. This feature relies on the flexibility of P's mechanism: for example, if we were to restrict P to only choose among stationary policies (i.e., those with a single offer α and a single threshold at which to make the offer), the choice of an optimal policy would depend on X_0.

6 Equilibrium

To evaluate the strength our commitment assumption, it is useful to think of where commitment comes from. In some situations, we think of commitment as coming via repeated interactions; for example, a firm interacting with many workers or, after agreeing to a split, having multiple interactions with the worker. Here, commitment based on the history of X_t may be possible if we punish deviations via movement to equilibria with low payoffs for P. In other situations, we think of commitment as coming from the ability to write court-enforceable contracts. In these contexts, our commitment assumption appears quite strong: writing a legally enforceable contract that depends on A's *past* outside option may be difficult to implement. Without such strong contracts, is implementing our mechanism-design solution possible? To answer this, we need to think carefully about the appropriate amount of commitment power to give P.

Much of the bargaining literature has focused on cases in which the offer that a player makes is "short lived" and expires at the end of a period. This inability to commit to offers past today has been shown to restrict the set of equilibria, often in very sharp ways (e.g., the Coasian bargaining literature). Such a stark lack of commitment is relevant in some bargaining frameworks; for example, in trade negotiations between countries, finding an enforcement body that allows players commit to offers may be hard.

However, we often see some forms of commitment used in real-world bargaining. Firms may be able to present a contract to a worker that the worker has time to consider before deciding whether to accept or reject. Thus, even if P doesn't have access to full commitment power, he may have access to some simpler forms of commitment that are easier for a contract to specify and a court to enforce.

In this section, we show even without full commitment power, we provide two ways in which we are able to approximate our mechanism-design solution. We start by designing a simple set of contracts which require only a small amount of commitment and possess a unique stationary equilibrium. We also proceed to show how to construct an equilibrium without *any* form of commitment in an alternating-offers framework which approximates our mechanism design solution. These results show that our main results still apply even when we relax the assumption of commitment.

6.1 Option with Escape Clause

Our first exercise looks at how to find a *simple* set of contracts that cannot explicitly condition on the outside option but will still approximate our mechanism design outcome. Although other contracts can potentially achieve this desiderata, we define a very simple one that we call an *option with escape clause*. This contract gives A an option with the right to, at any time, request a split with a prespecified share $1-\alpha$ of the pie. When A exercises this option, P can either accept the split or exercise the escape clause, which gives P the right to cancel the split by paying a penalty p. After paying the penalty, he can propose a new contract. We parameterize these contracts by the pair (α, p).

To understand the role of the escape clause, remember the optimal mechanism uses a split threshold $S(M_t)$ strictly above what A would choose. If P were to simply offer an option without an escape clause, A's resulting strategy would feature too much delay and A would take the split when X_t is too low relative to the optimal mechanism. The escape clause to the contract effectively allows P to select the threshold at which A will choose to exercise the option. If A's outside option is too low when he exercises the option, P will find it profitable to pay the penalty and propose a new option with a higher demand. P can use the penalty p to commit himself to accept the split if and only if A's outside option is sufficiently high. By setting the appropriate penalty, P will be able to ensure A takes the split at $S(M_t)$.

We study a discrete time game[11] with periods $t = 0, \Delta, 2\Delta, ...$ P can make long-lived offers in the form of an option with an escape clause or short-lived offers that expire at the end of the period. At the beginning of each period both players observe X_t and A, if holding an option, is given the choice to opt out or keep the option. If A opts out, P is given the chance to make a new offer to A. If A is holding an option, he is given the choice to exercise the option, take his outside option, or delay; if A is holding a short-lived offer, than A can accept the offer, reject the offer, or take his outside option. If A exercises the option, then P is given a choice to accept the split or pay a penalty p to use the escape clause. If either P cancels the split or A delays, then we move to the next period, discounting payoffs by $e^{-r\Delta}$.

In this new discrete-time game, the payoffs for A are the same as before, but the payoffs for P may change depending on whether P chooses to pay the escape penalty p. Let p_i the escape punishment in the ith option canceled and $\tau_{c,i}$ be the stopping time when P cancels the split for for ith time. P's payoff is given by

$$\mathbb{E}[e^{-r\tau} d_\tau u_P(\alpha_\tau) - \sum_{i=0}^{\infty} e^{-r\tau_{c,i}} p_i | X_0].$$

Because A is indifferent between opting out of a contract and P canceling the contract, we focus on equilibrium in which P never cancels the contract on-path (we can make the choice to opt out rather than force P to cancel strictly optimal for A by including a small transaction cost to exercise the option).

To simplify the proof, we assume X_t follows a random walk on a grid of points on $[\underline{X}, \overline{X}]$, moving up by one grid point with probability $q(X_t)$ and down by one grid point with probability $1 - q(X_t)$. When taking the limit as $\Delta \to 0$, we also take the distance of the grid points to 0 at an appropriate speed so that the random walk converges in distribution to our continuous-time diffusion process.[12] Analogously to our continuous-time assumptions, we assume $e^{-rt}X_t$ is still a strict super-martingale and that $q(X_t) \in (0, 1)\ \forall X_t \in (\underline{X}, \overline{X})$. Finally, we assume $\underline{X}$ is either natural[13] or absorbing.

[11]Using a discrete-time structure both fits with much of the bargaining literature and allows us to avoid well-known complications with equilibrium definition in continuous time.

[12]For example, the discrete-time random walk, when the grid size is properly scaled, is known to converge in distribution to a Brownian motion as $\Delta \to 0$. Daley and Green (2012) construct a discrete time random walk that converges to the belief distribution about a seller's type. Convergence to general processes of the form $dX_t = \mu(X_t)dt + \sigma dB_t$ (where $\mu(\cdot)$ is Hölder continuous) is shown in Gruber and Schweizer (2006).

[13]A boundary is *natural* if $\underline{X}$ cannot be reached in finite time. For example, if X_t is a discrete time version of a Geometric Brownian motion, it will be natural.

For this game, we focus on *stationary* equilibria. A stationary equilibrium requires P's equilibrium offers to be the same after histories h_t, h_s such that $X_t = X_s$ and A's equilibrium actions at t depend only the current value of X_t and the option or offer he holds. If a player ever deviates from their equilibrium offer, their opponent expects the deviation to be one-shot and play to return to the equilibrium path. Our main result shows that there exists a unique stationary subgame-perfect equilibrium outcome,[14] and that, as $\Delta \to 0$, P and A's equilibrium payoffs converge to the P-optimal mechanism payoffs of $J^*(X_0; X_0)$ and X_0, respectively.

Proposition 4. *For each Δ and given generic u_P, u_A and grid of points for X_t, every stationary subgame-perfect equilibrium leads to the same outcome. This outcome converges to our mechanism design solution with $W = X_0$ as $\Delta \to 0$.*

The literature in changing bargaining environments has focused on stationary strategies and short-lived offers. In our environment, as noted in Section 4, a stationary strategy using only short-lived offers would be unable to replicate the mechanism-design outcome. The option with an escape clause creates a "persistence" in P's offers and allows us to implement different split amounts and thresholds depending on when the option with escape clause is proposed. These contracts represent a simple way to retain stationarity while achieving a higher payoff for P.

6.2 Alternating-Offers

Proposition 4 tells us that with only a *small* amount of commitment, getting close to our mechanism-design solution is possible. However, as discussed earlier, settings exist in which even these simple long-term offers are not feasible. The majority of the bargaining literature has focused on such cases. A natural question is then what the best equilibrium outcome P could achieve when long-term offers such as the option with an escape clause are not available.

To answer this question, we study a canonical alternating-offers bargaining version of our model in which players make offers in a prespecified alternating order at $t = 0, \Delta, 2\Delta, ...$ Within each period, both players observe the realization of X_t, after which one player i is called to make a demand. Player k is then given a chance to either accept i's demand, reject i's demand or take his outside option. If k accepts i's demand, the game ends and the agreed upon split is made (with k getting the remaining share of the pie). If k rejects i's demand, then player i is given a chance to take his outside option

[14] The uniqueness result is only over outcomes, rather than strategies, because for any $\Delta > 0$, a small interval (whose size goes to 0 as $\Delta \to 0$) of penalties exists that can implement the same outcome.

or move to the next period. If i decides to move to the next period, then both players incur discounting costs $e^{-r\Delta}$. We will assume that the distribution of $X_{t+\Delta}$ conditional on X_t is the same as in the Section 3.

The optimal mechanism relies on P making credible promises to lower his demands in the future. We may be worried about whether P will fulfill these promises in equilibrium. When X_t is low and P is called to make an offer, he may be tempted to renege and increase his demand, knowing A's outside option is now lower. Foreseeing this possibility when at a high X_t today, A may not view P's promise to decrease his demand in the future as credible, and chooses to take his outside option today.

We prevent this kind of unravelling by specifying that when player i increases his demand higher than he is called to, we move to a punishment equilibrium in which their opponent k rejects the offer and makes a high demand in the next period. If k can credibly threaten to take the outside option upon i rejecting k's high demand, i will find it optimal to accept this high demand. This allows us to threaten i with a harsh punishment for increasing his demand. In the proof, we build subsequent off-path equilibria that make k's threat to the outside option credible.

Proposition 5. *Fix any $W \in (X_0, V^*(X_0))$. There exists a sequence of subgame-perfect equilibria as $\Delta \to 0$ with equilibrium payoffs J^Δ and W^Δ to P and A respectively such that $\lim_{\Delta \to 0} J^\Delta = J^*(X_0; W)$ and $\lim_{\Delta \to 0} W^\Delta = W$.*

Our proof constructs equilibria using strategies that approximate those we derived in the optimal mechanism and, therefore, maintain the same dynamics as in the optimal mechanism. If we allow P to choose his preferred equilibrium, as is standard in mechanism design, the loss to P from relaxing commitment is negligible for small Δ. Our analysis of these dynamics from the continuous-time case also then readily applies, telling us the on-path play of these equilibria stays close to the Pareto-frontier at all times.

7 Two-Sided Outside Option

So far, we have considered the case in which the outside option for P is equal to zero. However, in many situations, this will not be the case; in firm-worker negotiations, the firm's outside option may be to search for another worker. Additionally, the outside option for P and A may be correlated. For example, consider a firm bargaining with a union: the union's outside option is to go to strike. If we interpret X_t as the probability the strike will be successful, the outside options of the two players have a negative relationship. Similarly, we might think of a prosecutor bargaining with a defendant, where X_t is the probability the jury rules in favor of the defendant. In this case, the prosecutor may become more likely to strike a deal the higher X_t goes in order to avoid the risk that the defendant is acquitted in a trial.

We can formally introduce a richer outside option for P into the model by letting $v_P(X_t)$ be P's outside option. P's payoff can then be written as

$$\mathbb{E}[e^{-r\tau}\big(d_\tau(u_P(\alpha_\tau) - v_P(X_t)) + v_P(X_t)\big)|X_0].$$

Perhaps the first functional form that comes to mind is to assume P's outside option is a constant $v_P(X_t) = \nu$. This form is easily incorporated into our model without qualitatively changing any results. However, in the case where $v_P(X_t)$ is not a constant, we need to place some assumptions on the structure of v_P to make the analysis tractable. Assumption 3 ensures P's outside option isn't so large that he prefers the outside option to the best possible split that is individually rational for A, $\alpha = 1 - u_A^{-1}(X_t)$.

Assumption 3. *P's outside option satisfies $v_P(X_t) \leq u_P(1 - min\{u_A^{-1}(X_t), 1\})$.*

Assumption 3 preserves the following property from our baseline model with $v_P(X_t) = 0$: the best possible split for P which A will accept is better for P than his outside option. If X_t is too high, the best offer that P can make is to offer the entire pie to A. Using either this assumption or assuming v_P is constant, we can extend our optimal mechanism to allow for a richer set of outside options for P.

Theorem 2. *If v_P satisfies Assumption 3 or is a constant, there exists $\tilde{R} \leq \overline{R}$ and decreasing continuous functions $S(\cdot), \alpha(\cdot)$ such that optimal mechanism with two-sided outside options is given by $(\tau^*, \alpha^*_\tau, d_\tau)$, where*

$$\begin{aligned}
\tau^* &= inf\{t : \ X_t \geq \tilde{R} \text{ or } X_t \leq S(M_t)\} \\
d^*_\tau &= \mathbb{1}(X_\tau = S(M_\tau)) \\
\alpha^*_\tau &= \alpha(M_\tau).
\end{aligned}$$

The only real substantive difference comes from the fact that P may take the outside option earlier than in the case with $v_P(X_t) = 0$ (i.e., $\tilde{R} \leq \overline{R}$). The proof of Theorem 2 differs from that of Theorem 1 only in that we cannot directly apply to Lemma 1 to pin down the form of d_τ. Assumption 3 is used to pin down structure of the optimal decision rule. It allows us to rule out cases in which P takes the outside option at a low X_t and splits the pie at a higher X_t. Such a mechanism might be optimal if $v_P(X_t)$ decreased in X_t sufficiently fast. Assumption 3 lets us rule such a possibility out by ensuring that if P finds it optimal to make a split with A rather than take the outside option at X', then he will find it optimal to make a split at all $X < X'$ as well.

Throughout our analysis, our main focus in on preventing A from taking the outside option early. One concern in this extension is that P's outside option is positive, he might also be tempted to take the outside option early. Let us consider the simple case when P's outside option is a constant $\nu \geq 0$ and let $J^*_\nu(X_0; X_0)$ be P's payoff from the optimal mechanism when $W = X_0$ and P's outside option is ν. Although we rule out P taking the outside option early by our commitment assumption, this concern is important when we want to relax this assumption. Fortunately, we can extend the results of Proposition 5 to show alternating-offers equilibrium which approximates P's optimal mechanism exists.

Proposition 6. *In the discrete time alternating offers bargaining game in which P's outside option is $\nu \geq 0$, there exists a sequence of equilibria as $\Delta \to 0$ that deliver values $J^\Delta_\nu(X_0; X_0)$ to P such that $\lim_{\Delta \to 0} J^\Delta_\nu(X_0; X_0) = J^*_\nu(X_0; X_0)$.*

8 Conclusion

In this paper, we study a bargaining game in which one player's outside option may change over time. We find the outside option leads to a rich set of dynamics in the optimal bargaining outcome when one side can commit to their offers. The committed party gradually decreases the demands he makes and the pressure being placed on the other party over the course of the game, with periods of intransigence followed by quick spurts of concession reminiscent of haggling. Our model shows a new interplay between demands and pressure and finds they are complementary in providing incentives to continue bargaining. We characterize the Pareto frontier of the game and find similar dynamics arise in every point along the frontier.

In the second part of the paper, we explore how to relax the assumption that one party can commit to his offer process. We provide a simple set of contracts, namely options with escape clauses, that allow us to implement the mechanism-design solution as a unique stationary equilibrium. These contracts do not rely on outside courts to observe outside options or the history of play and can be easily enforced. Relaxing commitment even further, we study a classic alternating-offers bargaining game and find subgame perfect equilibrium which mimic our optimal mechanism when the period length becomes small. These exercises show us the dynamics generated by our optimal mechanism are robust to relaxing the commitment assumption.

References

Dilip Abreu and Faruk Gul. Bargaining and reputation. *Econometrica*, 68:85–117, 2000.

Thomas Balzer and Klaus Janßen. A duality approach to problems of combined stopping and deciding under constraints. *Mathematical Methods of Operations Research*, 55: 431–446, 2002.

Avner Shaked Binmore, Ken and John Sutton. An outside option experiment. *The Quarterly Journal of Economics*, 104:753–770, 1989.

Simon Board and Marek Pycia. Outside options and the failure of the coase conjecture. *American Economic Review*, 104:656–671, 2014.

Isaıas N Chaves. Privacy in bargaining: The case of endogenous entry. Working Paper, 2019.

Olivier Compte and Philippe Jehiel. On the role of outside options in bargaining with obstinate parties.. *Econometrica*, 70:1477–1517, 2002.

Martin W Cripps. Markov bargaining games. *Journal of Economic Dynamics and Control*, 22:341–355, 1998.

Brendan Daley and Brett Green. Waiting for news in the market for lemons. *Econometrica*, 80:1433–1504, 2012.

Brendan Daley and Brett S. Green. Bargaining and news. Working Paper, 2018.

Savas Dayanik and Ioannis Karatzas. On the optimal stopping problem for one-dimensional diffusions. *tochastic processes and their applications*, 107:173–212, 2003.

N.G. Dokuchaev. Optimal stopping of stochastic processes in a problem with constraints. *Th. Probab. Appl*, 41:761–768, 1997.

William Fuchs and Andrzej Skrzypacz. Bargaining with arrival of new traders. *American Economic Review*, 100:802–836, 2010.

David K. Levine Fudenberg, Drew and Jean Tirole. Infinite-horizon models of bargaining with one-sided incomplete information. *Game-theoretic Models of Bargaining (Alvin E. Roth, ed.)*, page 73–98, 1985.

Urs Gruber and Martin Schweizer. A diffusion limit for generalized correlated random walks. *Journal of applied probability*, pages 60–73, 2006.

Faruk Gul and Hugo Sonnenschein. On delay in bargaining with one-sided uncertainty. *Econometrica*, pages 601–611, 1988.

Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *Review of Economic Studies*, 49:315–333, 1982.

Ilwoo Hwang and Fei Li. Transparency of outside options in bargaining. *Journal of Economic Theory*, 167:116–147, 2017.

Jihong Lee and Qingmin Liu. Gambling reputation: Repeated bargaining with outside options. *Econometrica*, 81:1601–1672, 2013.

Andrew McClellan. Experimenation and approval mechanisms. Working Paper, 2019.

Antonio Merlo and Charles Wilson. A stochastic model of sequential bargaining with complete information. *Econometrica*, pages 371–399, 1995.

John Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

Juan Ortner. Durable goods monopoly with stochastic costs. *Theoretical Economics*, 12: 817–861, 2017.

Debraj Ray. The time structure of self-enforcing agreements. *Econometrica*, pages 547–582, 2002.

Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, pages 97–109, 1982.

Yuliy Sannikov. A continuous-time version of the principal-agent problem. *Review of Economic Studies*, pages 957–984, 2008.

Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *The Review of Economic Studies*, pages 541–554, 1988.

Appendix A

A.1 Lemma 1

Before stating the proof of Lemma 1, we show that there are sufficiently high X such that at X the optimal decision for A, if we were offered the entire pie, would be to stop and take the outside option immediately.

Lemma A.1. *The solution to $V^*(X)$ is, for some $\underline{S}, \overline{R}$, $\tau = \inf\{t : X_t \notin (\underline{S}, \overline{R})\}$ and $d_\tau = \mathbb{1}(u_A(1) \geq X_\tau$. $V^*(X) = X$ if and only if $X \geq \overline{R}$.*

Proof. Because $d_\tau = 1$ if and only $u_A(1) \geq X_\tau$, we can rewrite $V^*(X)$ as $\sup_\tau \mathbb{E}[e^{-r\tau} max\{u_A(1), X_\tau\}|X]$. $V^*(X)$ is bounded by the following observation:

$$\begin{aligned}\mathbb{E}[e^{-r\tau} max\{u_A(1), X_\tau\}|X] &= \mathbb{E}[e^{-r\tau} X_\tau|X] + \mathbb{E}[e^{-r\tau}(u_A(1) - X_\tau)\mathbb{1}(X_\tau < u_A(1))|X] \\ &\leq \mathbb{E}[e^{-r\tau} X_\tau|X] + u_A(1) \\ &\leq X + u_A(1).\end{aligned}$$

By standard optimal stopping arguments (see Dayanik and Karatzas (2003)) we can partition $[\underline{X}, \overline{X}]$ into an open continuation region $\mathcal{C} = \{V^*(X) > max\{u_A(1), X\}\}$ and a closed stopping region $\mathcal{D} = \{V^*(X) = max\{u_A(1), X\}\}$. The optimal stopping rule is $\tau = inf\{t : X_t \notin \mathcal{C}\}$.

We first argue that $\mathcal{C} = (\underline{S}, \overline{R})$ for some $(\underline{S}, \overline{R})$. Starting at $X_0 \in \mathcal{C}$, let $\tau_1 = inf\{t : X_t \in \mathcal{D},\ X_t \geq X_0\}$ be the first time X_t reaches $\mathcal{D}$ above X_0 and $\tau_2 = inf\{t : X_t \in \mathcal{D},\ X_t \leq X_0\}$ be the first time X_t reaches a $\mathcal{D}$ below X_0. By definition of the optimal τ, $V^*(X_0) = \mathbb{E}[e^{-r(\tau_1 \wedge \tau_2)} max\{u_A(1), X_{\tau_1 \wedge \tau_2}\}|X_0]$. By Assumption 2 and $V^*(X_0) > X_0$ by definition of $\mathcal{C}$, it must be that $X_{\tau_1} < u_A(1) < X_{\tau_2}$: If $X_{\tau_2} < u_A(1)$, then $V^*(X) = \mathbb{E}[e^{-r(\tau_1 \wedge \tau_2)} X_{\tau_1 \wedge \tau_2}|X_0] < X_0$ by Assumption 2; if $X_{\tau_1} > u_A(1)$, then $V^*(X) = \mathbb{E}[e^{-r(\tau_1 \wedge \tau_2)} u_A(1)|X_0] < u_A(1)$. Thus, for any $X_0 \in \mathcal{C}$, we must have, if $X_0 \leq u_A(1)$, $(X_0, u_A(1)] \subset \mathcal{C}$ or, if $X > u_A(1)$, $[u_A(1), X_0) \subset \mathcal{C}$. Taking $\underline{S} = X_{\tau_1}$ and $\overline{R} = X_{\tau_2}$, we conclude that $\mathcal{C} = (\underline{S}, \overline{R})$.

By definition of $\mathcal{C}$, we know that $V^*(X) = X$ only if $\mathcal{D}$. $\underline{S} < u_A(1) < \overline{R}$ implies that $V^*(X) = u_A(1)$ if $X \leq \underline{S}$ and $V^*(X) = X$ if $X \geq \overline{R}$, completing the proof.

□

A.1.1 Proof of Lemma 1

Proof. The if direction follows directly from the definition of $\overline{R}$: because, for $X > \overline{R}$, A can always achieve his first-best payoff $V^*(X)$ by quitting immediately, the only DIR

mechanism at is to take the outside option immediately. For the only if direction, suppose the optimal mechanism lets A take the outside option at $X_{t^*} < \overline{R}$. By definition of $\overline{R}$, if P were to thereafter offer the entire pie to A and let A choose when to take the entire pie, A would be induced to delay taking the outside option. For a small $\epsilon > 0$, P could decrease his demand slightly to $\alpha = \epsilon$ (again letting A decide when to take the offer) and A would still choose to delay taking the outside option (which we call an ϵ-offer continuation mechanism). It is straightforward to see that A would prefer to continue bargaining for ϵ sufficiently small. Let W_ϵ be A continuation value from this offer. By definition of our ϵ-offer mechanism, it must be that $W_\epsilon > X_\epsilon$.

Consider replacing $d_\tau = 0$ at τ such that $X_\tau < \overline{R}$ with our ϵ-offer continuation mechanism which induces A to delay taking his outside option. In the continuation game, A never finds it optimal to stop early because the mechanism thresholds are chosen to maximize A's utility. This new mechanism satisfies DIR because moving from X_τ to W_ϵ when $X_\tau < \overline{R}$ increases A's utility of following P's mechanism by

$$\mathbb{E}[e^{-r\tau}(W_\epsilon - X_\tau)\mathbb{1}(X_\tau < \overline{R} \textit{ and } \ d_\tau = 0)|X_0]$$

which is greater than the increase in A's utility when A deviates by taking the outside option at τ'

$$\mathbb{E}[e^{-r(\tau\wedge\tau')}\mathbb{1}(\tau < \tau')(W_\epsilon - X_\tau)\mathbb{1}(X_\tau < \overline{R} \textit{ and } \ d_\tau = 0)|X_0].$$

Moreover, this new mechanism is strictly better for P. Therefore choosing to stop at X_τ cannot have been optimal.

□

A.2 Lemma A.2

Before going to the proof of Proposition 2, we provide a useful characterization of the optimal stopping rule prior to $\tau_+(X^1)$. We show two important features of the optimal stopping rule, namely, that the stopping rule is locally stationary before $\tau_+(X^1)$ and the continuation mechanism at $\tau_+(X^1)$ is independent of $h_{\tau_+(X^1)}$.

Lemma A.2. *The solution to 2 is given by*

$$\begin{aligned}\tau &= \tau_-(S_0)\mathbb{1}(\tau_-(S_1) < \tau_+(X^1)) + (\tau_1 + \tau_+(X^1))\mathbb{1}(\tau_-(S_0) < \tau_+(X^1))\\ \alpha_\tau &= \alpha_0\mathbb{1}(\tau_-(S_0) < \tau_+(X^1)) + \alpha^1_\tau\mathbb{1}(\tau_-(S_0) < \tau_+(X^1))\end{aligned}$$

for some $(\alpha_0, S_0) \in \mathbb{R}^2$ and continuation mechanism (τ_1, α^1_τ) which is the same for all histories $h_{\tau_+(X^1)}$.

Proof. If the policy stops before $\tau_+(X^1)$, then the split α_τ solves

$$\max_{\alpha\in[0,1]} \; u_P(\alpha) - \eta u_A(1-\alpha). \tag{6}$$

The α which solves this is independent of the history of play. Let α_0 be the argmax of equation 6, which is unique by our concavity assumption.

Define a function $K(X^1)$ be the continuation value upon $\tau_+(X^1)$:

$$\begin{aligned} K(X^1) := \sup_{(\tau,\alpha_\tau)} \; \mathbb{E}\Big[& e^{-r\tau}(d_\tau u_P(\alpha_\tau) - (\eta + \lambda(X^1))\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big) \\ & + \sum_{n=2}^{N} \lambda(X_n)\big\{e^{-r(\tau\wedge\tau_+(X_n))}\big(d_\tau(X_n)\big(u_A(1-\alpha_\tau) - X_{\tau\wedge\tau_+(X_n)}\big) + X_{\tau\wedge\tau_+(X_n)})\big)] \\ & - e^{-r\tau}(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau)\big\}|X^1\Big] + \lambda(X^1)X^1 \end{aligned}$$

Applying the principal of optimality, we know the solution the continuation mechanism for (τ, α) upon reaching $\tau_+(X^1)$ must solve $K(X^1)$ and the value $K(X^1)$ is the same for all histories prior to $\tau_+(X^1)$. Let (τ^1, α^1_τ) be the solution to $K(X^1)$. The solution to 2 must solve

$$\begin{aligned} L(X_0) = \sup_{(\tau,\alpha_\tau)} \; \mathbb{E}\big[& e^{-r(\tau\wedge\tau_+(X^1))}\big(\big\{d_\tau[u_P(\alpha_\tau) - \eta(u_A(1-\alpha_\tau) - X_\tau)] \\ & - \eta X_\tau\big\}\mathbb{1}(\tau < \tau_+(X^1)) + K(X^1)\mathbb{1}(\tau \geq \tau_+(X^1))\big)|X_0\big]. \end{aligned}$$

We can divide $(\underline{X}, X^1)$ into an open continuation region $\mathcal{C}$ and a closed stopping region $\mathcal{D}$ defined by

$$\begin{aligned} \mathcal{C} &= \{X : \; L(X) > u_P(\alpha_0) - \eta u_A(1-\alpha_0)\} \\ \mathcal{D} &= \{X : \; L(X) = u_P(\alpha_0) - \eta u_A(1-\alpha_0)\}. \end{aligned}$$

Using Proposition 5.7 in Dayanik and Karatzas (2003), the optimal stopping rule τ is such that $\tau \wedge \tau_+(X^1) = \inf\{t : \; X_t \in \mathcal{D}\} \wedge \tau_+(X^1)$. First, we argue that $\mathcal{D}$ is not empty. If it were, then $V(\tau \wedge \tau_+(X^1), d_\tau(X^1), \alpha_\tau, X_0) = V(\tau_+(X^1), 0, \alpha_\tau, X_0) < X_0$. PK will be violated since $W \geq X_0$ and $RDIC(X^1)$ binding implies

$$X_0 > V(\tau \wedge \tau_+(X^1), d_\tau(X^1), \alpha_\tau, X_0) = V(\tau, d_\tau, \alpha_\tau, X_0) \geq W \geq X_0.$$

Finally, we claim that $\mathcal{D}$ must take the form of a connected interval $[\underline{X}, S_0]$ for some S_0. Suppose this were not the case. Then $\exists D_1, D_2 \in \mathcal{D}$ such that $X \in \mathcal{C}$ for all $X \in (D_1, D_2)$.

For such X we have

$$L(X) = \mathbb{E}[e^{-r(\tau_-(D_1)\wedge\tau_+(D_2))}\big(L(D_1)\mathbb{1}(\tau_-(D_1) < \tau_+(D_2)) \\ + L(D_2)\mathbb{1}(\tau_-(D_1) > \tau_+(D_2))\big)|X_0 = X]$$

Because α_0 is fixed, we have $L(D_1) = L(D_2)$. Therefore $L(X)$ is strictly less than the utility from making the split with α_0 immediately. This is contradicts $X \in \mathcal{C}$.

□

A.3 Lemmas A.3 and A.4

Here we prove a couple of useful facts about the continuation values of A in the optimal mechanism for our relaxed problem that are used in the proof of Proposition 2 to show how the optimal threshold and S, α change with M_t. The arguments here use the local stationarity of the solution to our relaxed problem as well as complementary slackness conditions to provide some useful properties of A's continuation value in the relaxed problem.

Lemma A.3. *For each $X^m \in \mathcal{B}$, A's continuation value at $\tau_+(X^m)$ is X^m.*

Proof. Let $(\tau[X^m], d_\tau[X^m], \alpha[X^m])$ be the continuation mechanism at $\tau_+(X^m)$. By the arguments in Proposition 2, this continuation mechanism will be the same for all histories prior to $\tau_+(X^m)$. The fact that the $RDIR(X^m)$ constraint binds implies that

$$\mathbb{E}[e^{-r(\tau\wedge\tau_+(X^m))}\big(d_\tau(X^m)(u_A(1-\alpha_\tau) - X_\tau) + X_{\tau\wedge\tau(X^m)}\big)|X_0] \\ = \mathbb{E}[e^{-r\tau}\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0].$$

We can rewrite each side of the equation splitting apart the events $\tau < \tau_+(X^m)$ and $\tau \geq \tau_+(X^m)$. The deviation payoff is

$$\mathbb{E}[e^{-r\tau}\mathbb{1}(\tau < \tau_+(X^m))\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0] \\ + \mathbb{E}[e^{-r\tau_+(X^m)}\mathbb{1}(\tau \geq \tau_+(X^m))X^m|X_0],$$

while the payoff from following the mechanism is

$$\mathbb{E}[e^{-r\tau}\mathbb{1}(\tau < \tau_+(X^m))\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0] \\ + \mathbb{E}[e^{-r\tau}\mathbb{1}(\tau \geq \tau_+(X^m))\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0].$$

Together these imply that

$$
\begin{aligned}
&\mathbb{E}[e^{-r\tau_+(X^m)}\mathbb{1}(\tau \geq \tau_+(X^m))X^m|X_0] \\
&= \mathbb{E}[e^{-r\tau}\mathbb{1}(\tau \geq \tau_+(X^m))\big(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau\big)|X_0] \\
&= \mathbb{E}[e^{-r\tau_+(X^m)}\mathbb{1}(\tau \geq \tau_+(X^m))\mathbb{E}[e^{-r\tau[X^m]}\big(d_\tau[X^m](u_A(1-\alpha[X^m]) - X_{\tau[X^m]}) + X_{\tau[X^m]}\big)|X^m]|X_0].
\end{aligned}
$$

Because the optimal policy at $\tau_+(X^m)$ is independent of the previous history of play, we can treat $\mathbb{E}[e^{-r\tau[X^m]}\big(d_\tau[X^m](u_A(1-\alpha[X^m]) - X_{\tau[X^m]}) + X_{\tau[X^m]}\big)|X^m]$ as a constant. Pulling it out of the expectation, we have

$$
X^m = \mathbb{E}[e^{-r\tau[X^m]}\big(d_\tau[X^m](u_A(1-\alpha[X^m]) - X_{\tau[X^m]}) + X_{\tau[X^m]}\big)|X^m].
$$

Therefore, the policy $(\tau[X^m], d_\tau[X^m], \alpha[X^m])$ yields a continuation value of X^m for A. □

Lemma A.4. *For any $M \in (X^m, X^{m+1})$, A's continuation value in the optimal relaxed mechanism at (X, M) is below X if and only if $X \in (X^m, X^{m+1})$.*

Proof. By Lemma A.3 we know that A's continuation value at $\tau_+(X^m)$ is equal to X^m and at $\tau_+(X^{m+1})$ is equal to X^{m+1}. Take an arbitrary (X, M) such that $X \in (X^m, X^{m+1})$ and $M < X^{m+1}$. Because the mechanism is stationary until $\tau_+(X^{m+1})$, A's continuation value will continue to be X^m at all (X^m, M) with $M < X^{m+1}$. This implies that we can express A's continuation value at (X, M) as

$$
\begin{aligned}
&\mathbb{E}[e^{-r(\tau_-(X^m)\wedge\tau_+(X^{m+1}))}\{X^m\mathbb{1}(\tau_-(X^m) < \tau_+(X^{m+1})) + X^{m+1}\mathbb{1}(\tau_-(X^m) > \tau_+(X^{m+1}))\}|X] \\
&< X,
\end{aligned}
$$

where the inequality follows from Assumption 2 and Doob's Optional Stopping Theorem.

To show the only if direction, for the sake contradiction let $X < X^m$ and suppose $W \leq X$ was A's continuation value at (X, M). By Lemma A.3, we know A's continuation value at $\tau_+(X^{m+1})$ is X^{m+1}. Again using the stationarity of the optimal mechanism until $\tau_+(X^{m+1})$, A's continuation value at $\tau_+(X^m)$ is

$$
\begin{aligned}
&\mathbb{E}[e^{-r(\tau_-(X)\wedge\tau_+(X^{m+1}))}(\mathbb{1}(\tau_-(X) < \tau_+(X^{m+1}))W + \mathbb{1}(\tau_-(X) > \tau_+(X^{m+1}))X^{m+1}|X^m] \\
&\leq \mathbb{E}[e^{-r(\tau_-(X)\wedge\tau_+(X^{m+1}))}(\mathbb{1}(\tau_-(X) < \tau_+(X^{m+1}))X + \mathbb{1}(\tau_-(X) > \tau_+(X^{m+1}))X^{m+1}|X^m] \\
&< X^m,
\end{aligned}
$$

contradicting the fact that A's continuation value was X^m. Therefore, $W > X$. □

A.4 Proof of Proposition 2

Proof. The structure of τ, α follows directly from a repeated application of Lemma A.2 as discussed in the text. To see that α_m is decreasing in m, consider a maximization problem of the form

$$\max_{\alpha \in [0,1]} \; u_P(\alpha) - \gamma u_A(1-\alpha),$$

where $\gamma = \eta + \sum_{k=1}^{m} \lambda(X^k)$ (if $m = 0$ we take $\gamma = \eta$). The optimal choice of α is increasing in γ. We can conclude that α_m is decreasing in m because $\lambda(X^k) < 0$ implies that γ is decreasing in m.

We prove that S_m is decreasing in m through the following series of observations. Our first, discussed in the text, notes that P will always set the threshold S_m so that A takes the split before A would optimally choose to.

Observation 1. *The optimal threshold S_m must be above where A would optimally choose to stop and take a split of α_m.*

We will look at the choice of the optimal threshold at a fixed X before and after X^m has been reached. Take (X, M^j), $j = m-1, m$ with $X \in [X^{m-1}, X^m)$, $M^{m-1} \in [X^{m-1}, X^m)$ and $M^m \in [X^m, X^{m+1})$. Define $H(X^m)$ be the value to P of the continuation mechanism at $\tau_+(X^m)$.[15] By Lemma A.3 we know that A's continuation value at $\tau_+(X^m)$ is equal to X^m. Because the optimal mechanism is stationary between $\tau_+(X^m)$ and $\tau_+(X^{m+1})$, P and A's continuation values at X^m will be $H(X^m)$ and X^m, respectively, regardless of whether they are starting at (X, M^{m-1}) or (X, M^m).

Our next observation notes that A's continuation value will be lower at (X, M^{m-1}) than at (X, M^m). The intuition for this observation is simple: the higher M has been, the more P must increase A's continuation value to incentivize A to not take his outside option early. By Lemma A.4, we know that $W^{m-1} < X < W^m$.

Observation 2. *The promised continuation value to A at (X, M) is increasing in M: $W^m > W^{m-1}$.*

We know there will be a static threshold S at which P will implement a split. In Lemma E.14 we show that this threshold is unique. Given W and a threshold S, we can pin down the split amount $\alpha(W, S)$ needed to deliver the utility W to A:

$$\alpha(S, W) = 1 - u_A^{-1}\left(\frac{W - \Psi(X^m, S, X)X^m}{\psi(X^m, S, X)}\right).$$

[15] By the independence of the continuation mechanism with respect to $h_{\tau_+(X^m)}$, the P's continuation value will be the same at every realization of $\tau_+(X^m)$.

We can then define P's problem as $\max_S\ F(S,W)$ where

$$F(S,W) := \Psi(X^m, S, X)H(X^m) + \psi(X^m, S, X)u_P(\alpha(W,S)),$$

which has a first-order condition of

$$\frac{\partial F(S,W)}{\partial S} = \Psi_S H(X^m) + \psi_S u_P(\alpha) + \psi u_P'(\alpha)\alpha_S = 0.$$

Because P benefits from raising S (i.e., $\Psi_S H(X^m) + \psi_S u_P(\alpha) > 0$), it must be that $\alpha_S < 0$ for the first-order condition to hold.

Let $S(W)$ be the optimal choice of S when delivering W utility to A. To know the sign of $S'(W)$, we need to see whether the cross-partial term $\frac{\partial F(S,W)}{\partial W \partial S}$ is positive or negative. Taking the derivative respect to W, we have

$$\frac{\partial F(S,W)}{\partial W \partial S} = \psi_S u_P'(\alpha)\alpha_W + \psi u_P''(\alpha(W,S))\alpha_W \alpha_S + \psi u_P'(\alpha)\alpha_{WS}. \tag{7}$$

Using the functional form of $\alpha(S,W)$, we get that

$$\psi u_P'(\alpha)\alpha_{WS} = -\psi_S u_P'(\alpha)\alpha_W - \alpha_S \frac{u_A''(1-\alpha)u_P'(\alpha)}{u_A'(1-\alpha)}\psi.$$

Plugging this into equation 7 and using the fact that both $\alpha_W, \alpha_S < 0$, we get $\frac{\partial F(W)}{\partial W \partial S} < 0$. Using the second-order condition on F, we have that

$$S'(W) = \frac{-\frac{\partial^2 F(S(W),W)}{\partial W^2}}{\frac{\partial F(S(W),W)}{\partial W \partial S}} < 0.$$

Finally, we argue that the choice of $S(W^1)$, $S(W^2)$ are the optimal thresholds in Proposition 2. Suppose that P instead used another threshold which delivered utility W^i to A and utility $\tilde{j}(X)$ to P at X while using $S(W^i)$ delivered utility $j(X)$ to P. By the choice of $S(W^i)$, we know $j(X) \geq \tilde{j}(X)$. Then for every $X' > X$, we could replace the continuation value at X with $j(X)$, increasing P's utility at X' while still giving the same continuation value to A at every X'. But because we know that a single threshold is optimal, this means that the optimal threshold X' is the same as at X. Therefore it must be that the optimal mechanism at X uses threshold $S(W^i)$.

□

A.5 Proof of Theorem 1

Proof. The solution to our relaxed problem is given by Proposition 2. Let us look at the limit as the grid $\mathcal{X}_N$ becomes arbitrarily fine: $\mathcal{X}_N = \{X_n : X_n = X_0 + (i-1)\frac{\overline{R}-X_0}{N} \; for \; i \in \{1, ..., N\}\}$. For each N, let $S_N(M)$ and $\alpha_N(M)$ be the functions describing the optimal split threshold and split amount in the relaxed problem using $\mathcal{X}_N$. Our limit mechanism can then be defined as the limit of $S_N(M), \alpha_N(M)$. Because S_N, α_N are monotonic and bounded, the limit is well-defined by Helly's Selection Theorem. Using the continuity of P's payoffs with respect to the choice of S, α, it follows that this limit mechanism provides an upper-bound on P's problem.

To show that the limit mechanism solves our full problem, we verify that A's continuation value is weakly greater than his outside option everywhere. Because the mechanism is measurable with respect to (X_t, M_t), A's continuation value will also be measurable with respect to (X_t, M_t) as well. By Lemma A.3, we know that when $X_t = M_t$, A's continuation value from following the mechanism is equal to X_t. Suppose that at some $X_t = X' < M_t$, A's continuation value was strictly less than X'.

By Lemma A.4 we know that for each N, the continuation value for A in our relaxed problem is only strictly below his outside option only when $X_t \in (X^m, X^{m+1})$ and $M_t \in (X^m, X^{m+1})$. This implies that in our relaxed problem there is no grid point between (X^m, X^{m+1}); if there was, then $RDIR$ would be violated on this grid point. As we take our grid to be finer and finer, we get that the distance between X^m, X^{m+1} goes to zero. Because $X' < M_t$ and A's payoffs are continuous with respect to α, S, A's continuation value at (X', M_t) must be strictly negative in the solution to the relaxed problem for all sufficiently large N. But this can only happen if X' is between M_t and the previous grid point in $\mathcal{X}_N$ below M_t. This implies that $X' \in (M_t - \frac{\overline{R}-X_0}{N}, M_t)$, which cannot be for large enough N. Therefore no such X' can exist. We conclude that the limit mechanism has a weakly positive continuation value for all histories, thereby satisfying DIR. Finally, continuity of α, S is shown in Lemma A.5 below.

□

A.6 Proof of Continuity of S, α

Lemma A.5. *Both $\alpha(\cdot), S(\cdot)$ are continuous functions.*

Proof. Consider a stopping problem in which A must choose when to take his outside option when he receives the split $u_A(1-\alpha)$ at $\tau_-(S)$ (where α, S remain fixed). By dynamic consistency, he will choose to take the outside option at some threshold $B(\alpha, S)$ which is independent of the starting X. By the same arguments as in Lemma A.2 we know this threshold will be unique and, using the Theorem of the Maximum, will be continuous in α, S.

We argue that, in the optimal limit mechanism, the mechanism only adjusts (i.e., changes α, S) when A would optimally choose to take his outside option if the mechanism were to remain fixed. Because A chooses to optimally take the outside option at $B(\alpha, S)$ and we know that the mechanism only adjusts at $X = M$, this is equivalent to $B(\alpha(M), S(M)) = M$. For the sake of contradiction suppose that for some M_1 and $\epsilon > 0$ we had $B(\alpha(M_1), S(M_1)) = M_1 + \epsilon$.[16] Define a new mechanism $(\widehat{\alpha}(M), \widehat{S}(M))$ which is identical to $(\alpha(M), S(M))$ on $M \notin (M_1, M_1 + \epsilon)$ but keeps $\alpha(M), S(M)$ fixed over $[M_1, M_1 + \epsilon]$ (so that $\widehat{\alpha}(M) \geq \alpha(M)$ and $\widehat{S}(M) \geq S(M)$). We will argue that this higher demand and threshold lead to a strict increase in P's utility. That this new mechanism satisfies our DIR constraints is clear; we know that A's continuation value on $(M, M + \epsilon)$ under $\widehat{\alpha}, \widehat{S}$ greater than the outside option by definition of $B(\alpha, S)$. A's continuation value at M is strictly higher than before since, under the original mechanism, α, S only change when A's DIR constraint is binding.

Let τ^* be the stopping rule in our original mechanism and $\widehat{\tau}$ be the stopping rule in our modified mechanism. To show that this new mechanism is better for P, we will split the possible paths $X_t[\omega]$ of sample point ω into two cases: those which lead to A taking outside option and those which lead to a split. This first case is given by $\{\omega : \ \tau_+(\overline{R})[\omega] = \tau^*\}$. Taking expectation over such ω, using $(\widehat{\alpha}, \widehat{S})$ strictly increases P's payoffs because P now reaches a split on the set $\{\omega : \ inf\{t : \ X_t[\omega] \leq \widehat{S}(M[\omega])\} < \tau_+(\overline{R})[\omega] = \tau^*\}$ (which gives higher utility than the outside option) and leads to the same payoff on $\{\omega : \ inf\{t : \ X_t[\omega] \leq \widehat{S}(M[\omega])\} \geq \tau_+(\overline{R})[\omega] = \tau^*\}$. Therefore, on $\{\omega : \ \tau_+(\overline{R})[\omega] = \tau^*\}$, P prefers $(\widehat{\alpha}, \widehat{S})$ to (α, S).

Next, consider the set $\{\omega : \ \tau^* \neq \tau_+(\overline{R})[\omega]\}$. Again using $(\widehat{\alpha}, \widehat{S})$ weakly increases P's utility over such ω because

$$\widehat{\tau} := inf\{t : X_t \leq \widehat{S}(M_t)\} \wedge \tau_+(\overline{R}) \leq inf\{t : X_t \leq S(M_t)\} \wedge \tau_+(\overline{R}) = \tau^*,$$

which follows from the fact that $S(M)$ is decreasing in M. Therefore, because P gets an earlier split (decreasing discounting costs) and receives a higher terminal split $\alpha_{\widehat{\tau}} \geq \alpha_{\tau^*}$ (which follows from the facts that $\alpha(M)$ is decreasing in M, $\widehat{\alpha}(M) \geq \alpha(M)$ and $\widehat{\tau} \leq \tau^*$).

Putting these observations together together, P is strictly better off under $(\widehat{\alpha}, \widehat{S})$, contradicting the optimality of (α, S). Therefore it must be that $B(\alpha(M), S(M)) = M$. This then implies that (α, S) must be continuous. If either had a jump, then $B(\alpha, S)$ would also jump, which cannot be if $B(\alpha(M), S(M)) = M$.

[16] It cannot be that $B(\alpha(M), S(M)) < M$, as this would lead to a violation of A's DIR constraint between $B(\alpha(M), S(M))$ and M.

□

A.7 Proof of Proposition 3

Proof. Suppose that after some history h_t continuation value for both P, A was off the Pareto frontier. Let W_t be the continuation value for A after such a history. By replacing the continuation mechanism after h_t with an optimal mechanism with a $PK(W_t)$ constraint we could strictly increase P's payoffs. Moreover, this wouldn't change the incentive of A to take his outside option before t since A evaluates the continuation value after h_t in the new mechanism as the same as in the old mechanism. Therefore A's continuation value at every history h_s which might lead to h_t is exactly the same: if A had no strict incentive to take his outside option at s in the old mechanism, then he will have no incentive to take his outside option at s in the new mechanism.

□

Appendix B

B.1 Options with Escape Clause

B.1.1 Notation

Take an arbitrary stationary equilibrium. Let $J^{\Delta}(X, \alpha, p)$ and $V^{\Delta}(X, \alpha, p)$ be the equilibrium value functions for P and A respectively when the current state is X and the current contract is (α, p); a short-term offer of α will be given by $(\alpha, \emptyset)$. We define $J^{\Delta}(X) := J^{\Delta}(X, \alpha_X, p_X)$ where (α_X, p_X) is the equilibrium contract offered at X and, simiarily for A, $V^{\Delta}(X) := V^{\Delta}(X, \alpha_X, p_X)$.

We can write out A's equilibrium value function as

$$V^{\Delta}(X, \alpha, p) = max\{\sup_{\tau, d_\tau} \mathbb{E}[e^{-r\tau}\big(d_\tau \mathbb{1}\big(J^{\Delta}(X_\tau) - p \leq u_P(\alpha)\big)u_A(1-\alpha) \\ + (1 - d_\tau \mathbb{1}\big(J^{\Delta}(X_\tau) - p \leq u_P(\alpha)\big))V^{\Delta}(X_\tau, \alpha_{X_\tau}, p_{X_\tau})\big)|X], X\},$$

where, slightly abusing notation, $d_\tau \in \{0, 1\}$ is equal to 0 if A opts out of the contract and 1 if A exercises the option. We say that A agrees to a split if either he accepts a short-lived offer by P or he exercises the option and P agrees to not cancel the split. Note that we are already imposing that P cancels the contract whenever $J^{\Delta}(X_\tau) - p > u_P(\alpha)$, a necessary condition in any stationary subgame perfect equilibrium. By standard optimal stopping arguments, A's optimal stopping strategy consists of $\tau = inf\{t : \ X_t \notin (b_X, B_X)\}$ for some thresholds b_X, B_X (for $(\alpha, \emptyset)$, if A accepts immeidately then $b_X = B_X = X$;

otherwise, if A rejects, $b_X = X - \epsilon,\ B_X = X + \epsilon$). We will let $(\tau(\alpha, p), d_\tau(\alpha, p))$ be A's optimal strategy when the current option is (α, p).

Let $C^*(X)$ be the set of contracts (α, p) such that A doesn't immediately find it optimal to take the outside option when P offers (α, p) at X. P's value function from the stationary equilibrium when proposing a new offer is then

$$J^\Delta(X) = \sup_{(\alpha,p)\in C^*(X)} \mathbb{E}[e^{-r\tau(\alpha,p)}\big(max\{u_P(\alpha), J^\Delta(X_{\tau(\alpha,p)}) - p\}d_\tau(\alpha,p)) \\ + (1 - d_\tau(\alpha,p))J^\Delta(X_{\tau(\alpha,p)})\big)|X]. \tag{8}$$

If $C^*(X)$ is empty, then A must take the outside option immediately and $J^\Delta(X) = 0$. When the current contract is (α, p), P's value function is

$$J^\Delta(X,\alpha,p) = \mathbb{E}[e^{-r\tau(\alpha,p)}\big(max\{u_P(\alpha), J^\Delta(X_{\tau(\alpha,p)}) - p\}d_\tau(\alpha,p) + (1 - d_\tau(\alpha,p))J^\Delta(X_{\tau(\alpha,p)})\big)|X].$$

We will define Ψ^Δ, ψ^Δ to be the same expected discounted probabilities as before but now accounting for the discrete time law of X. Finally, for each grid point on which X_t can fall, we can define $\epsilon^+_{X_t}$ to be the distance upward and $\epsilon^-_{X_t}$ to be the distance downward to the next grid point when starting at X_t. To simplify notation, we will drop let ϵ denote both $\epsilon^+_{X_t}$ and $\epsilon^-_{X_t}$. None of the arguments rely on the exact distance $\epsilon^+_{X_t}$ and $\epsilon^-_{X_t}$ so this abuse of notation should not cause confusion.

B.2 Supporting Lemmas

We first must prove a several supporting Lemmas that will prove useful in the equilibrium characterization. Our first Lemma makes a simple observation about the value functions for A, P at X where A doesn't choose to opt out of a contract (α, p).

Lemma B.6. *If A doesn't opt out of an offer (α, p) at X, then $V^\Delta(X, \alpha, p) \geq V^\Delta(X, \alpha_X, p_X)$ and $J^\Delta(X, \alpha_X, p_X) \geq J^\Delta(X, \alpha, p)$.*

Proof. The inequality for A is immediate from A's choice to not opt out. For P, it follows from the fact that he could choose to offer (α, p) at X and A would take the same actions as if he held the offer (α, p) at the start of the period. Hence P must do at least as well with the offer (α_X, p_X). □

The next Lemma shows that as $X \to \underline{X}$, A's value function V^Δ must be equal to his outside option and a split will be agreed to immediately.

Lemma B.7. *For all sufficiently low X, we have $V^\Delta(X) = X$ and A agrees to a split immediately.*

Proof. For the sake contradiction, suppose that for any X', we can find an $X \in [\underline{X}, X')$ such that $V^\Delta(X) > X$. As $X \to \underline{X}$, the discounted probability that A takes the outside option must also go to zero.[17] Let τ_e be the equilibrium time until a split is reached and $\alpha_e = \mathbb{E}[\alpha_{\tau_e}|X]$. Then $V^\Delta(X) \approx \mathbb{E}[e^{-r\tau_e} u_A(1-\alpha_{\tau_e})|X]$ and $J^\Delta(X) \approx \mathbb{E}[e^{-r\tau_e} u_P(\alpha_{\tau_e})|X]$. If the equilibrium doesn't call for A to immediately accept, we know by the concavity of u_P, u_A,

$$u_P(\alpha_e) > \mathbb{E}[e^{-r\tau_e} u_P(\alpha_{\tau_e})|X]$$
$$u_A(1-\alpha_e) > \mathbb{E}[e^{-r\tau_e} u_A(1-\alpha_{\tau_e})|X].$$

P could make a short-lived offer of α_e and improve both player's utilities if taken immediately. Therefore, for all sufficiently low X, A accepts the equilibrium offer immediately.

Next, we argue that P has a profitable deviation if $V^\Delta(X) > X$. Suppose that P changes the equilibrium offer from α_X to $\alpha' = \alpha_X + \delta$. If A accepts immediately, he receives $u_A(1 - \alpha_X - \delta)$ while if he opts out and takes the next-period split he gets

$$e^{-r\Delta}[q(X)u_A(1-\alpha_{X+\epsilon}) + (1-q(X))u_A(1-\alpha_{X-\epsilon})].$$

Because A accepts immediately, it must be that A is indifferent between accepting immediately and waiting until the next period (otherwise P could increase his demand and still induce A to accept immediately). But, because this holds at every X for which A is called to agree to a split immediately, A must be indifferent between taking an offer today and waiting until the equilibrium calls for him not to accept immediately. Because the delay in this alternate strategy becomes arbitrarily long as $X \to \underline{X}$, we have a contradiction unless $\alpha_X \to 1$. But this cannot be since $V^\Delta(X) \geq \underline{X} > u_A(0)$.

□

In the next Lemma, we argue that if, at some X, the equilibrium leads to an immediate split, then it does so at all lower X. In the optimal mechanism this property is intuitive: the lower X is, the lower the value of experimentation. Lemma B.8 shows that the same holds in equilibrium.

Lemma B.8. *If A immediately agrees to a split with a new offer $(\alpha_{X'}, p_{X'})$ at X', then A also does so for all (α_X, p_X) at $X < X'$.*

Proof. For the sake of contradiction, suppose there exist X_1, X_2 such that A agrees to a split immediately at X_1, X_2 and all $X < X_1$ but doesn't accept immediately at any

[17] If A were called to take the outside option at sufficiently low X, P would have a profitable deviation to offer $\alpha = 1 - u_A^{-1}(X) - \delta$ for some sufficiently small δ, inducing A to not take the outside option.

$X \in (X_1, X_2)$. Consider the case in which, at some $X \in (X_1, X_2)$, A will accept a split with probability one in the future. Consider an alternative offer P could make at X in which P makes a short-lived offer $(\alpha_e, \emptyset)$ where τ_e, α_e are as defined in Lemma B.7. This increases A's utility since $V^\Delta(X) = \mathbb{E}[e^{-r\tau_e} u_A(1-\alpha_{\tau_e})|X] < \mathbb{E}[u_A(1-\alpha_{\tau_e})|X] \leq u_A(1-\alpha_e)$ and will be accepted immediately. This also increases P's utility for a similar reason and therefore represents a profitable deviation.

For A to not accept a split with probability one, it must be that for some X A's best response to P's equilibrium offer involves A opting out of the split at $B_X > X_2$.[18] We note that $V^\Delta(X_2, \alpha_X, p_X) \geq V^\Delta(X_2, \alpha_{X_2}, p_{X_2})$ because A always has the option of opting out of the contract at X_2. The inequality also goes in the opposite direction; at X_2 with offer (α_X, p_X), we know that A's continuation value is equal to

$$V^\Delta(X_2, \alpha_X, p_X) = \Psi^\Delta(B_X, X, X_2)V^\Delta(B_X, \alpha_{B_X}, p_{B_X}) + \psi^\Delta(B_X, X, X_2)V^\Delta(X, \alpha_X, p_X).$$

A has a deviation strategy in which he can hold on to the contract at X_2 until either B_X or X is reached and then opt out. For this to not be a profitable deviation by A, we must have

$$\begin{aligned} V^\Delta(X_2, \alpha_{X_2}, p_{X_2}) &\geq \Psi^\Delta(B_X, X, X_2)V^\Delta(B_X, \alpha_{B_X}, p_{B_X}) + \psi^\Delta(B_X, X, X_2)V^\Delta(X, \alpha_X, p_X) \\ &= V^\Delta(X_2, \alpha_X, p_X). \end{aligned}$$

Therefore A's utility is the same as if he were to opt out of (α_X, p_X) at X_2, in which case he would accept immediately; that is, there is a strategy that ends with accepting a split with probability one that is optimal for A. By Lemma B.6 we know that $J(X_2, \alpha_{X_2}, p_{X_2}) \geq J(X_2, \alpha_X, p_X)$ and so P would be weakly better off if A were to opt out at X_2. This new strategy which always leads to split would be better for both players. Defining α_e with respect to this "always split" strategy, P can make an offer of $(\alpha_e, \emptyset)$ that leaves both strictly better off.

□

The next Lemma establishes that P's value function is decreasing in A's outside option. When J^Δ is decreasing, the choice of p is effectively the choice of a threshold S below which P will always use the escape clause and above which he will not. This property will allow P to select the threshold at which A will exercise the escape clause.

Lemma B.9. *$J^\Delta(X)$ is decreasing in X.*

[18] Opting out of a split at $b_X < X_1$ would lead to an immediate split by definition of X_1.

Proof. For all X sufficiently large, A will choose to take his outside option and $J^{\Delta}(X) = 0$. This implies $J^{\Delta}(X)$ must be weakly decreasing somewhere. If it is not weakly decreasing everywhere, there is an X such that $max\{J^{\Delta}(X-\epsilon), J^{\Delta}(X+\epsilon)\} < J^{\Delta}(X)$.

First, suppose that A doesn't immediately agree to a split at X. If A doesn't opt out of the offer (α_X, p_X) at either $X-\epsilon$ or $X+\epsilon$, then we have

$$\begin{aligned} J^{\Delta}(X) &= e^{-r\Delta}\Big[q(X)J^{\Delta}(X+\epsilon, \alpha_X, p_X) + (1-q(X))J^{\Delta}(X-\epsilon, \alpha_X, p_X)\Big] \\ &\le e^{-r\Delta}\Big[q(X)J^{\Delta}(X+\epsilon) + (1-q(X))J^{\Delta}(X-\epsilon)\Big] \\ &< max\{J^{\Delta}(X-\epsilon), J^{\Delta}(X+\epsilon)\}, \end{aligned}$$

a contradiction. If A opts out only at $X-\epsilon$. Then we have

$$\begin{aligned} J^{\Delta}(X) &= e^{-r\Delta}\Big[q(X)J^{\Delta}(X+\epsilon, \alpha_X, p_X) + (1-q(X))J^{\Delta}(X-\epsilon)\Big] \\ &\le e^{-r\Delta}\Big[q(X)J^{\Delta}(X+\epsilon) + (1-q(X))J^{\Delta}(X-\epsilon)\Big] \\ &< max\{J^{\Delta}(X-\epsilon), J^{\Delta}(X+\epsilon)\}, \end{aligned}$$

a contradiction. A similar argument applies if A opts out at only $X+\epsilon$ or both $X+\epsilon$ and $X-\epsilon$.

The only possibility left is that A agrees to a split immediately at X. Suppose P deviates at $X-\epsilon$ and makes an offer of α_X. If A were to accept this, this would yield a strict improvement for P. The only way A would opt out at $X-\epsilon$ is if A's continuation value was higher than $u_A(1-\alpha_X)$.

But, by Lemma B.8, we know A must accept immediately at $X-\epsilon$. Moreover, it must be that A is indifferent between accepting immediately and delaying (otherwise P could increase his demand and A would still find it optimal to accept the split immediately). Therefore, we know

$$u_A(1-\alpha_{X-\epsilon}) = e^{-r\Delta}[q(X-\epsilon)u_A(1-\alpha_X) + (1-q(X-\epsilon))u_A(1-\alpha_{X-2\epsilon})]. \quad (9)$$

If $u_A(1-\alpha_{X-\epsilon}) > u_A(1-\alpha_X)$, equation 9 implies $u_A(1-\alpha_{X-2\epsilon}) > u_A(1-\alpha_{X-\epsilon})$. Applying the same arguments at $X-2\epsilon$, P's demand is even smaller at $X-3\epsilon$ and so on. But for a low enough outside option, this will violate Lemma B.7. Therefore it must be that A would choose to accept α_X at $X-\epsilon$.

□

The following Lemma concerns A's preferences over thresholds at which to implement a split. Fix some some X and $b < X$. Let $\widehat{V}$ be A's utility from various stopping thresholds $B \geq X$ when his continuation value at b is fixed to be $V^{\Delta}(b)$:

$$\widehat{V}(B) = \Psi^{\Delta}(B, b, X)u_A(1-\alpha) + \psi^{\Delta}(B, b, X)V^{\Delta}(b).$$

Lemma B.10. *$\widehat{V}(B)$ is single-peaked in $B \geq X$ and $B^* := \max_{B \geq X} \widehat{V}(B)$ is increasing α.*

Proof. For the sake of contradiction, suppose that $\widehat{V}(B)$ was not single-peaked. Then there are $B_1 < B_2 < B_3$ such that $\widehat{V}(B_2) < min\{\widehat{V}(B_1), \widehat{V}(B_3)\}$. Suppose $V(B_1) \leq V(B_3)$ and define ρ so that A's utility is the same from stopping at B_1 and B_3 if stopping at B_3 yields 0 utility with probability $1-\rho$ (the case when $V(B_1) > V(B_3)$ follows by the same argument when attaching ρ to stopping at B_1):

$$\begin{aligned}&\Psi^{\Delta}(B_1, b, X)u_A(1-\alpha) + \psi^{\Delta}(B_1, b, X)V^{\Delta}(b)\\&= \Psi^{\Delta}(B_3, b, X)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, X)V^{\Delta}(b).\end{aligned}$$

Define a new function $\widehat{V}_{\rho}(B)$ in which we modify the utility of stopping in $\widehat{V}(B)$ to include this ρ:

$$\widehat{V}_{\rho}(B) := \Psi^{\Delta}(B, b, X)u_A(1-\alpha)(1 + \mathbb{1}(B = B_3)(\rho - 1)) + \psi^{\Delta}(B, b, X)V^{\Delta}(b).$$

By definition of $\widehat{V}_{\rho}$, we have $\widehat{V}_{\rho}(B_1) = \widehat{V}_{\rho}(B_3)$. Using fact that the continuation value upon reaching B_1 is $\Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)$, we can write the utility from using B_3 as:

$$\widehat{V}_{\rho}(B_3) = \Psi^{\Delta}(B_1, b, X)\big[\Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)\big] + \psi^{\Delta}(B_1, b, X)V^{\Delta}(b).$$

The utility from using B_1 can be written as

$$\widehat{V}_{\rho}(B_1) = \Psi^{\Delta}(B_1, b, X)u_A(1-\alpha) + \psi^{\Delta}(B_1, b, X)V^{\Delta}(b).$$

Together these imply that $u_A(1-\alpha) = \Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)$. Using $\widehat{V}_{\rho}(B_3) > \widehat{V}_{\rho}(B_2)$, we have

$$\begin{aligned}\widehat{V}_{\rho}(B_3) &= \Psi^{\Delta}(B_2, b, X)\big[\Psi^{\Delta}(B_3, B_1, B_2)u_A(1-\alpha)\rho\\&\quad + \psi^{\Delta}(B_3, B_1, B_2)\big(\Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)\big)\big] + \psi^{\Delta}(B_2, b, X)V^{\Delta}(b)\\&> \Psi(B_2, b, X)u_A(1-\alpha) + \psi(B_2, b, X)V^{\Delta}(b)\\&= \widehat{V}_{\rho}(B_2).\end{aligned}$$

Simplifying this inequality, we get

$$\begin{aligned}&\Psi^{\Delta}(B_3, B_1, B_2)u_A(1-\alpha)\rho \\ &+ \psi^{\Delta}(B_3, B_1, B_2)\big(\Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)\big) \\ &> u_A(1-\alpha).\end{aligned} \tag{10}$$

Because $u_A(1-\alpha) = \Psi^{\Delta}(B_3, b, B_1)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, b, B_1)V^{\Delta}(b)$, the left side of the inequality in equation 10 is equal to

$$\Psi^{\Delta}(B_3, B_1, B_2)u_A(1-\alpha)\rho + \psi^{\Delta}(B_3, B_1, B_2)u_A(1-\alpha) < u_A(1-\alpha),$$

a contradiction. Therefore no such B_1, B_2, B_3 exist and $\widehat{V}$ must be single-peaked.

Finally, we argue that B^* is increasing in α. If A strictly prefers one choice of B to all others, then this will continue to hold for all small changes in α. Consider the point at which A is indifferent between two choices $B' < B''$:

$$\Psi^{\Delta}(B', b, X)u_A(1-\alpha) + \psi^{\Delta}(B', b, X)V^{\Delta}(b) = \Psi^{\Delta}(B'', b, X)u_A(1-\alpha) + \psi^{\Delta}(B'', b, X)V^{\Delta}(b)$$

Increasing α will decrease the side with the larger Ψ^{Δ} the most. Because Ψ^{Δ} is decreasing in B, this will imply that A strictly prefers B''. □

An analogous result holds if we flip the roles of B and b. Taking the continuation value to be fixed at $B > X$ and considering A's utility over choices of b, we will conclude that A's utility is single-peaked in b and the optimal choice of b is *decreasing* in α; as with the choice of B^*, this means that the continuation region grows. We would need to switch Ψ^{Δ} and ψ^{Δ} but otherwise the proof follows directly from the arguments above.

Our main Lemma for showing equilibrium uniqueness establishes that every equilibrium offer (α_X, p_X) leaves A indifferent between waiting and taking his outside option.

Lemma B.11. *In every stationary equilibrium* $V^{\Delta}(X) = X \ \forall X$.

Proof. For every contract (α, p) that A doesn't opt out immediately at X, by Lemma B.6, $V^{\Delta}(X, \alpha, p) \geq V^{\Delta}(X, \alpha_X, p_X)$ and $J^{\Delta}(X, \alpha_X, p_X) \geq J^{\Delta}(X, \alpha, p)$. For any contract with which both hold with equality, let us assume that A to opts out at X. This keeps the same equilibrium payoffs for both players and will remain an equilibrium.

Let X be the smallest X' such that $V^{\Delta}(X') > X'$ (if there exists any X' such that $V(X') > X'$, then a smallest such X' exists by Lemma B.7). We will consider a deviation by P when proposing a new offer at $X_t = X$ to $(\alpha', p') = (\alpha_X + \delta, p_X + u_P(\alpha_X + \delta) - u_P(\alpha_X))$ (with $p' = \emptyset$ if $p_X = \emptyset$). This choice of a deviation penalty is picked so that P

will make the same decisions on whether to cancel the option if A attempts exercise the split. For small δ, if A was willing to agree to a split at X' when holding (α_X, p_X) rather than opt out, then he will still agree to a split whenever $V^\Delta(X', \alpha, p) > V^\Delta(X', \alpha_{X'}, p_{X'})$. If $V^\Delta(X', \alpha, p) = V^\Delta(X', \alpha_{X'}, p_{X'})$, then by our assumption on the equilibrium, we must have $J^\Delta(X', \alpha_{X'}, p_{X'}) > J^\Delta(X', \alpha, p)$ in which case P would be strictly better off if A opts out. Let us therefore focus on the case when $V^\Delta(X', \alpha, p) > V^\Delta(X', \alpha_{X'}, p_{X'})$ at X' at which A chooses to exercise the option.

We know A chooses to stop at $inf\{s \geq t : X_s \notin (b_X, B_X)\}$ when the equilibrium offer is (α_X, p_X). Consider the case in which A accepts the split at both B_X and b_X. If he accepts the split immediately he will get a strictly higher utility than taking the split at both b_X and B_X as long as P does not use the escape clause. P will not do so since his value function at X is $\mathbb{E}[e^{-r(\tau_+(B_X)\wedge\tau_-(b_X))}u_P(\alpha_X)|X]$, which is strictly less than $u_P(\alpha_X)$. Therefore, if A doesn't accept immediately at X, then it must be that he opts out of the contract at either b_X or B_X

Let us turn to when A exercises the split only at $b_X < X$. If $X_{t+\Delta} = X + \epsilon$, A expects to be able to opt out the next time X is reached and receive the equilibrium contract (α_X, p_X); therefore his continuation utility above X will be the same and the deviation in the offer will not affect A's choice of B_X. Consider A's decision when $X_{t+\Delta} = X - \epsilon$. We know that, because $J^\Delta(X)$ is decreasing, the original equilibrium can be implemented using $p_X = J^\Delta(b_X) - u_P(\alpha_X)$ (P will accept the split if and only if the outside option is greater than or equal to b_X). Using this as p_X, we know from the observation after Lemma B.10 that increasing α will never cause A to change the threshold at which he would choose to take the split; increasing α leads to A preferring to take the split at a lower threshold than b_X. A knows that P will cancel the contract at any $X_s < b_X$ and the continuation contract will deliver A his outside option. Therefore A will still accept at b_X and P will be strictly better off.

Next, we consider when A exercises the split only at $B_X > X$.[19] If $X_{t+\Delta} = X - \epsilon$, A's utility will be the same because he will opt out the next time X is reached and so A will still have the same incentive to opt out of the contract at b_X. The only change to A's strategy may be his choice of B_X. By Lemma B.10, we know that A's utility is single-peaked in the choice of B. Let B_X^* be A's optimal choice of when to take the split when P cannot cancel contract and A has continuation value $V^\Delta(b_X)$ at b_X:

$$B_X^* = \max_B \ \Psi(B, b_X, X)u_A(1-\alpha) + \psi(B, b_X, X)V^\Delta(b_X).$$

[19]This type of strategy, with acceptance only at higher X, is qualitatively different than our optimal mechanism. Although we could show that such offers will never be used in equilibrium, we need the fact that $V^\Delta(X) = X$ to establish this; for now we must consider the use of such strategies.

Because $V^{\Delta}(b_X) = b_X$ and $u_A(1-\alpha) > B_X^* > b_X$, it must be that A would be strictly better off taking the split immediately-i.e., $B_X^* = X < B_X$. By the fact that A's preferences are single-peaked in B_X, $B_X^* < B_X$ implies that A's utility is decreasing in B at B_X. A small increase in α will not induce A to choose a larger B_X and the fact that he was not choosing a lower B implies that P would cancel any contract at $X_s \leq B_X$. Therefore (α', p') is a strict improvement for P, as it doesn't change when A would accept or opt out of the contract and with some probability P gets a larger share of the split when A exercises the option.

In the case where A is called to agree to a split immediately at X, we know A must be indifferent between taking the split today or opting out of the contract and taking a new contract tomorrow; if this were not the case, then P could increase α and A would still find it optimal to accept immediately. We are left with the conclusion that if $V^{\Delta}(X) > X$, it must be either A accepts the offer immediately at X or A opts out at both B_X and b_X. Opting out at X is never strictly optimal. Instead of opting out at X, A could wait to opt out at B_X, b_X. Because this will lead to the same equilibrium outcomes, we will therefore restrict attention to strategies by A in which he never opts out at such X.

Next, consider the equilibrium contract at $X + \epsilon$. We can repeat almost all of our same arguments as above. The only argument we need to consider is the case where A accepts at B_X and A opts out at b_X with $V^{\Delta}(b_X) > b_X$. But by our previous argument we know that when $V^{\Delta}(b_X) > b_X$ and A opts out at b_X, it must be that A accepts the equilibrium offer at b_X immediately. When A is called to accept the offer immediately at b_X, we can improve on the contract by having P make a short-lived offer $\alpha' = \Psi^{\Delta}(B_X, b_X, X)\alpha_X + \psi^{\Delta}(B_X, b_X, X)\alpha_{b_X}$. This will improve both A and P's utility. Therefore, $V^{\Delta}(X+\epsilon) > X + \epsilon$ only if A is called to accept the contract immediately at $X + \epsilon$. We can repeat our arguments at $X + 2\epsilon, X + 3\epsilon$ and so on. We are left with the conclusion that for any X, if $V^{\Delta}(X) > X$ then A must be agreeing to a split immediately at X or opting out at both b_X, B_X.

Take the largest set $[X_1, X_2]$ (potentially with $X_1 = X_2$) with $V^{\Delta}(X) > X$ for all $X \in [X_1, X_2]$. A must be indifferent between taking the current offer and getting the equilibrium value tomorrow.[20] Delaying at each $X \in [X_1, X_2]$ will yield the same payoff for A as his equilibrium strategy. Using this, we can write A's utility at X as

$$V^{\Delta}(X) = \Psi^{\Delta}(X_2+\epsilon, X_1-\epsilon, X)V^{\Delta}(X_2+\epsilon) + \psi^{\Delta}(X_2+\epsilon, X_1-\epsilon, X)V^{\Delta}(X_1-\epsilon).$$

By definition of X_1, X_2, we know $V^{\Delta}(X_1-\epsilon) = X_1 - \epsilon$ and $V^{\Delta}(X_2+\epsilon) = X_2 + \epsilon$. By

[20]This been argued earlier when A accepts the split immediately. It is also true if A expects to opt out at both b_X, B_X.

Assumption 2, this means $V^{\Delta}(X) < X$, a contradiction. Therefore no such set $[X_1, X_2]$ can exist.

□

B.2.1 Proof of Proposition 4

Proof. First, we argue that a stationary equilibrium does indeed exist. We will let P choose a threshold S below which A cannot take the split (this is translated into a choice of p by setting $J^{\Delta}(S) = u_P(\alpha) - p$) and will replace p with S in all of our value functions. We know that upon opting out of the contract A expects to receive his outside option, making it easy to pin down his optimal strategy. Let A's value function from not taking an action today be

$$V_+^{\Delta}(X, \alpha, S) = \sup_{\tau \geq \Delta, d_\tau} \mathbb{E}[e^{-r\tau}(d_\tau \mathbb{1}(X_\tau \geq S)u_A(1-\alpha_\tau) + (1 - d_\tau \mathbb{1}(X_\tau \geq S))X_\tau | X].$$

Then we can write A's value function today as

$$V^{\Delta}(X, \alpha, S) = max\{V_+^{\Delta}(X, \alpha, S), X\}.$$

This leads to an optimal choice by A to stop at $b_{\alpha,S}$ and $B_{\alpha,S}$. Because of discreteness issues, A may be indifferent between two choices of b or B. In this case we allow P to choose his preferred threshold (this is without loss, since P could always perturb the equilibrium contract slightly to induce A to strictly prefer one threshold). We can recursively define P's value function as in equation 8. Let $C_S^*(X) = \{(\alpha, S) : \ V_+^{\Delta}(X, \alpha, S) \geq X\}$ be the set of contracts that don't induce A to take his outside option immediately. P's value function is

$$J^{\Delta}(X) = \sup_{(\alpha,S) \in C_S^*(X)} \mathbb{E}[e^{-r\tau(\alpha,S)}\big(d_\tau(\alpha,p)u_P(\alpha)) + (1 - d_\tau(\alpha,S))J^{\Delta}(X_\tau)\big)|X]$$

When A accepts at $b_{\alpha,S}$, this is equal to $\Psi^{\Delta}(B_{\alpha,S}, b_{\alpha,S}, X)J^{\Delta}(B_{\alpha,S}) + \psi^{\Delta}(B_{\alpha,S}, b_{\alpha,S}, X)u_P(\alpha)$ and is $\Psi^{\Delta}(B_{\alpha,S}, b_{\alpha,S}, X)u_P(\alpha) + \psi^{\Delta}(B_{\alpha,S}, b_{\alpha,S}, X)J^{\Delta}(b_{\alpha,S})$ when A accepts at $B_{\alpha,S}$. It is easily seen that $b_{\alpha,S}$ and $B_{\alpha,S}$ are bounded in α, S. Because A's utility is continuous in α and the set of thresholds S is finite, we know that we can replace the *sup* with *max* in $J^{\Delta}(X)$. Standard dynamic programming techniques then give us that $J^{\Delta}(X)$ exists and is unique. Taking the equilibrium actions to be (α_X, S_X) equal to the *argmax* of $J^{\Delta}(X)$, we have our equilibrium.

Next, we argue that there is a (generically) unique equilibrium choice of α, S. For each S_i, let α_i be the demand by P such that A is indifferent between taking the outside option at X and taking the contract. If P is indifferent between S_1, S_2 at X, we have

$$\Psi^{\Delta}(X+\epsilon, S_1, X)u_P(\alpha_1) + \psi^{\Delta}(X+\epsilon, S_1, X)J^{\Delta}(X+\epsilon)$$
$$= \Psi^{\Delta}(X+\epsilon, S_2, X)u_P(\alpha_2) + \psi^{\Delta}(X+\epsilon, S_2, X)J^{\Delta}(X+\epsilon)$$

We are using the fact that A will choose to opt out at $X+\epsilon$. If he chose to opt out of the contract at some $B_X > X+\epsilon$, then we know that his continuation value at X and B_X is equal to the outside option and hence his continuation value at $X' \in (X, B_X)$ would be strictly below X'.

Because A is indifferent between taking the contract and taking the outside option now, we know that

$$X = \Psi^{\Delta}(X+\epsilon, S_i, X)u_A(1-\alpha_i) + \psi^{\Delta}(X+\epsilon, S_i, X)(X+\epsilon).$$

Suppose we perturb the X grid by adding $\eta > 0$ to each point (keeping the probabilities of moving up or down the same). Then we need to adjust each α_i a bit to restore A's indifference. This implies that

$$\frac{d\alpha_i}{d\eta} = \frac{\psi^{\Delta}(X+\epsilon, S_i, X) - 1}{\Psi^{\Delta}(X+\epsilon, S_i, X)u_A'(1-\alpha_i)}$$

Using this changes P's utility from (α_i, S_i) by $\frac{(\psi^{\Delta}(X+\epsilon,S_i,X)-1)u_P'(\alpha_i)}{u_A'(1-\alpha_i)}$. For arbitrary u_P, u_A these will in general be different for $S_1 \neq S_2$ and P will strictly prefer one of the two. Therefore, if there were multiple (α_i, S_i) which were optimal before, one of them will be strictly optimal after this perturbation.

Finally, we argue that, as $\Delta \to 0$, the limit of P's equilibrium value converges to that of our full mechanism. It must be that the limit value of this discrete time game converges to something weakly below our continuous time limit, because we can always approximate the discrete time mechanism in continuous time. Therefore we only need to check that the discrete time equilibrium value doesn't converge to something strictly below the continuous time mechanism.

Because the discrete time random walk converges in distribution to the continuous time diffusion process, when the current contract is (α, S) we know A's optimal strategy will converge to the continuous time limit. Consider a (potentially sub-optimal) deviation for P in which he offers $(\alpha(M_t) - \delta, S(M_t) - \delta)$ at each X_t for some small δ (where $\alpha(M_t), S(M_t)$ come from the optimal continuous time mechanism). For small enough Δ, A would always choose to accept this option whenever P offers it. Because A's optimal choice of $b_{\alpha,S}, B_{\alpha,S}$ converges to his choice in the continuous time limit, by convergence of the distribution of the random walk to the continuous time diffusion process it then follows that P's value function will approach the continuous time limit with policy $(\alpha(M) - \delta, S(M) - \delta)$. We can take δ as small as we like and approximate the continuous time mechanism arbitrarily closely.

□

B.3 Proof of Proposition 5

Proof. Proposition 5 is a special case of Theorem 3 in Appendix D, which extends the model to allow P to have an outside option of $\nu \geq 0$. Theorem 3 requires that P's individual rationality constraint is not violated. This always holds when P's outside option is zero and so for any $W \in (X_0, V^*(X_0))$, we can use the equilibrium in Theorem 3 to approximate our optimal mechanism.

□

Appendix C

C.1 Proof of Theorem 2

Proof. The proof proceeds as in Theorem 1, defining the same relaxed problem with $RDIR$ constraints. Because we cannot directly apply Lemma 1, we must first pin down the structure of d_τ in the relaxed mechanism. The proof of Lemma A.2 goes through in an almost identical manner other than choice of d_τ.

Consider P's optimal policy at X_0 and consider the possible stopping rules used by P:

1. $d_\tau = 0$ at a lower threshold S_0, continue at $\tau_+(X^1)$.
2. $d_\tau = 0$ at a lower threshold S_0 and $d_\tau = 0$ at an upper threshold $R_0 \leq X^1$.
3. $d_\tau = 0$ at a lower threshold S_0 and $d_\tau = 1$ at an upper threshold $R_0 \leq X^1$.
4. $d_\tau = 1$ at a lower threshold S_0 and $d_\tau = 1$ at an upper threshold $R_0 \leq X^1$.
5. $d_\tau = 1$ at a lower threshold S_0 and $d_\tau = 0$ at a threshold $R_0 \leq X^1$.
6. $d_\tau = 1$ at a lower threshold S_0 and continue at $\tau_+(X^1)$

We can immediately rule out 1) and 2) as violating A's DIR constraints[21]. 4) we can rule out because the agreed upon split is independent of X_τ and P would be better off implementing the split immediately.

To rule out 3), we note that, if v_P is not constant, by Assumption 3 P would be better taking a split with $\alpha = 1 - u_A^{-1}(S_1)$ (if $1 - u_A^{-1}(S_1) > 0$) which leaves the agent indifferent between taking the outside option S_1 and taking the split. In order to not violate PK at X_0, it must be that the split amount α_1 at R_1 has $u_A(1-\alpha_1) > R_1$. Therefore, we know $1 - u_A^{-1}(R_1) > 0 \Rightarrow 1 - u_A^{-1}(S_1) > 0$. We conclude that it cannot be optimal to take the outside option at S_1. If $v_P(X_t) = \nu$ is constant, we know that $d_\tau = 1$ if

$$\max_\alpha u_P(\alpha) + \eta u_A(1-\alpha) > \nu + \eta X_\tau.$$

The fact that this is holds at R_1 means it holds at S_1, implying that it is optimal to take the split at S_1 and 4) cannot hold.

The only possible solutions are 5), 6). If we have a solution of the form 5), or applying the same argument for the continuation mechanism at $\tau_+(X^1)$, except for the structure of S_m thresholds, we can immediately conclude that solution to the relaxed problem has the same structure as in Proposition 2. The rest of the arguments in Theorem 1 then apply.

Finally, when $v_P(X_t) = \nu$, we prove the property that P's continuation value in the optimal mechanism at any (X_t, M_t) is strictly greater than his outside option except at $X_t = M_t = \tilde{R}$. This is obvious in the case where $\nu \leq 0$, because any delay in taking the outside option yields a continuation value to P strictly above ν. Therefore let us focus on the case when $\nu > 0$.

We know $J(X_0, X_0) > \nu$ and that $J(X_t, M_t)$ is strictly decreasing in M_t (this follows from the fact that the P share of the split and the split threshold are both decreasing in M_t). When, at some M_t, the DIR constraint is binding, it must be be that $J(M_t, M_t) \geq \nu$ (otherwise P could take the outside option at M_t and deliver the same utility to A and increase his own utility).

Suppose that $J(M_t, M_t) > \nu$ but for some small δ $J(X_t, M_t) < \nu - \delta$. Take some small ϵ_1 and $\epsilon_2 < \delta$. Using $J(M_t + \epsilon_1, M_t + \epsilon_1) \leq J(M_t, M_t)$ and $J(X_t, M_t') \leq J(X_t, M_t)$ for $M_t' \in (M_t', M_t + \epsilon_1)$, we can bound $J(M_t, M_t)$:

[21] Ruling out 1) follows from Lemma A.3 which tells us that A's continuation value at $\tau_+(X^1)$ is X^1.

$$
\begin{aligned}
J(M_t, M_t) &\leq \Psi(M_t + \epsilon_1, X_t, M_t)J(M_t, M_t) + \psi(M_t + \epsilon_1, X_t, M_t)(J(X_t, M_t) + \epsilon_2) \\
&\leq \Psi(M_t + \epsilon_1, X_t, M_t)J(M_t, M_t) + \psi(M_t + \epsilon_1, X_t, M_t)(\nu - \delta + \epsilon_2) \\
&< \Psi(M_t + \epsilon_1, X_t, M_t)J(M_t, M_t) + \psi(M_t + \epsilon_1, X_t, M_t)\nu \\
&< \Psi(M_t + \epsilon_1, X_t, M_t)J(M_t, M_t) + \psi(M_t + \epsilon_1, X_t, M_t)J(M_t, M_t) \leq J(M_t, M_t).
\end{aligned}
$$

a contradiction. Whenever at the first M_1 such that DIR is binding for A we have $J(M_1, M_1) > \nu$, then we have that $J(X_t, M_t)$ will be greater than ν for all (X_t, M_t).

□

Appendix D

D.1 Alternating Offers with P's outside option

Here we prove a more general statement than Propositions 5 and 6. The sacrifice we must make is that we require P's continuation value to be greater than ν after all histories, which we will henceforth assume. This assumption is too strong in many cases of interest: we know that this is true in P's optimal mechanism with $PK(X_0)$ and in the case for general W when $\nu = 0$.

D.1.1 Equilibrium Definition

We write P's demands at time t as α_t and A's demands at time t as β_t. Let $\alpha(M_t; W, X_0)$ $S(M_t; W, X_0)$, $\tilde{R}$ be the offer process and thresholds induced by the optimal mechanism starting at X_0 with $PK(W)$. We drop dependence on W, X_0 for notational convenience.

We generate our equilibrium using on-path strategies which approximate the optimal mechanism. Fix a small $\delta > 0$ and $\overline{R}^{\Delta} \approx \tilde{R}$ such that in $u_P(\alpha(\overline{R}^{\Delta}) - \delta)$ is bounded away from ν; as we take $\delta \to 0$, we will also take $\overline{R}^{\Delta} \to \widetilde{R}$. At each (X_t, M_t) such that $X_t > S(M_t)$, the proposing player demands the entire pie and the non-proposing player rejects. The first time $X_t \leq S(M_t)$, the proposing player demands a perturbed version of the optimal mechanism's demand: P demands $\alpha_t = \alpha(M_t) - \delta$ and A demands $\beta_t = 1 - \alpha(M_t) + \delta$. The non-proposing player is called to accept this offer. Both players immediately take the outside option if $X_t \geq \overline{R}^{\Delta}$.

In order to support these on-path strategies, we specify two types of punishment equilibrium. We define an *outside equilibrium*, which punishes a player by inducing all players to take their outside options whenever possible, and an *inside equilibrium*, which punishes a player by having their opponent raise their demands. The structure of the inside equilibrium will depend on the identity of player we are trying to punish.

For the inside equilibrium to punish A at X_t, we move to an equilibrium in which P uses an approximation of the optimal mechanism which delivers continuation value X_t to A. Let $M_{ts} := \max_{t' \in [t,s]} X_{t'}$ is the maximum of the process restarting at X_s. If the punishment equilibrium starts at time t, P uses the same strategies as if we were on path, only now using modifications as in our on-path strategies to the mechanism with $\alpha(M_{ts}; X_t, X_t), S(M_{ts}; X_t, X_t)$ rather than $\alpha(M_t; W, X_0), S(M_t; W, X_0)$. By, as in our path play, subtracting δ for P's demands, we ensure that A's DIR constraint will be satisfied in the discrete-time game. We then proceed as in the on-path equilibrium. A's continuation value, for small Δ, δ, is

$$\begin{aligned}\mathbb{E}[e^{-r\tau}u_A(1-\alpha_t)|X] &\approx \mathbb{E}[e^{-r\tau}(u_A(1-\alpha(M_{ts}; X_t, X_t)) + u'_A(1-\alpha(M_{ts}; X_t, X_t))\delta)|X] \\ &\approx X_t + \mathbb{E}[e^{-r\tau}u'_A(1-\alpha(M_{ts}; X_t, X_t))\delta|X] \\ &\leq X_t + k\delta,\end{aligned}$$

where $k = \max_{\beta} \; u'_A(\beta)$.

For the inside equilibrium to punish P at X_t, we specify an equilibrium in which A makes a high demand in the following period. In this equilibrium A, at $t+\Delta$, makes an offer of $\beta_{t+\Delta}$ which is picked so as to leave P indifferent between taking the outside option at t and waiting for A's offer at $t+\Delta$. Because $X_{t+\Delta}$ is stochastic, defining $\beta_{t+\Delta}$ requires a little care to ensure that A's demand is IR for A. For some δ_P to be determined shortly, let $\beta_I = 1 - u_P^{-1}(e^{r\Delta}\nu) - \delta_P$. In the inside equilibrium, A demands $\beta_{t+\Delta} = \beta_I$ if $X_{t+\Delta} + k\delta \leq u_A(\beta_I)$, demands $\beta_{t+\Delta} = 1 - u_P^{-1}(\nu)$ if $X_t \in (u_A(\beta_I) - k\delta, u_A^{-1}(1 - u_P^{-1}(\nu))]$ ($u_A^{-1}(1-u_P^{-1}(\nu))$ is maximal amount A can take demand which P will accept) and takes the outside option otherwise.[22] We choose the smallest δ_P which leaves P indifferent between taking his outside option at t or waiting until $t+\Delta$:

$$\nu = e^{-r\Delta}\big[u_P(1-\beta_I)\mathbb{P}(X_{t+\Delta} + k\delta \leq u_A(\beta_I)|X_t) + \nu\mathbb{P}(X_{t+\Delta} + k\delta > u_A(\beta_I)|X_t)\big].$$

If no such δ_P exists, then A demands $\beta_{t+\Delta} = 1 - u_P^{-1}(\nu)$ if $u_A(1 - u_P^{-1}(\nu)) \geq X_{t+\Delta}$ and take the outside option otherwise. When $X_{t+\Delta} + k\delta \leq u_A(\beta_I)$, P accepts A's demand if and only if $\beta_{t+\Delta} \leq \beta_I$. When $X_{t+\Delta} + k\delta > u_A(\beta_I)$, P accepts A's demand if and only if $\beta_{t+\Delta} \leq 1 - u_P^{-1}(\nu)$.

In the outside equilibrium each proposing player makes the maximal possible offer subject to their opponent's IR constraint (i.e., $\alpha_t = 1 - u_A^{-1}(X_t),\ \beta_t = 1 - u_P^{-1}(\nu)$). If

[22] By splitting A's demands according to whether or not $X_t + k\delta \gtrless u_A(\beta_I)$, we will be able to ensure that A's utility is high enough that he doesn't prefer to move to inside punishment equilibrium for A.

the other player rejects this offer, the proposing player immediately takes the outside option. All players are called to accept any offer that is individually rational.

We use these inside and outside equilibria to deter players from deviating in their demand/acceptance/rejection strategies by specifying the following reaction to a deviation. If player i makes a demand greater than he is called to when player k is called to accept the equilibrium offer and i's deviating offer is IR for k, then player k rejects the offer and we move to an inside equilibrium to punish player i. If player i makes a demand that is not IR when k is called to accept, then k takes the outside option and we move to an outside equilibrium. In any subgame (on- or off-path), if player i rejects an offer they are called to accept, then player k is called to take the outside option immediately and we move to the outside equilibrium for every subsequent subgame.

Before proving that this is an equilibrium, we show a useful property of P's discrete-time value function from this equilibrium. We prove that P's IR constraint holding in the continuous time limit mechanism means that, for small Δ, P's IR constraint will continue to hold when using our discrete time approximation described above.

Lemma D.12. *Let $J(X_t, M_t; X_0)$ be the continuation value for P at (X_t, M_t) from the optimal mechanism starting at X_0 in Theorem 2. Suppose that $J(X_t, M_t; X_0) > \nu$ for all $M_t < \tilde{R}$. For sufficiently small δ and $\overline{R}^{\Delta} < \tilde{R}$ but close to $\tilde{R}$, as $\Delta \to 0$, our discrete time equilibrium value function $J^{\Delta}(X_t, M_t)$ satisfies $J^{\Delta}(X_t, M_t) > \nu$ for all $X_t < M_t \leq \overline{R}^{\Delta}$.*

Proof. Let $\widehat{J}^{\Delta}(X, M)$ be the discrete-time continuation value to P at (X, M) when using the offer process $\alpha(M_t) - \delta$ for all $M \leq \tilde{R}$ (this is almost identical to the continuation value on-path in our equilibrium, only we are not restricting players to take the outside option before $\tilde{R}$). For some $\overline{R}^{\Delta} < \tilde{R}$ but close to $\tilde{R}$, by convergence of payoffs as $\Delta \to 0$, we know that for small δ, Δ, we will have $\widehat{J}^{\Delta}(X, M) > \nu$ for all $M > \overline{R}^{\Delta}$. We want to argue that if we force players to take the outside option at $\overline{R}^{\Delta}$, then the new equilibrium value function for P, $J^{\Delta}(X, M)$, will still have $J^{\Delta}(X, \overline{R}^{\Delta}) \geq \nu$ for all $X \leq \overline{R}^{\Delta}$. The utility at a particular $\widehat{J}^{\Delta}(X, M)$ for P is given by

$$\begin{aligned}\widehat{J}^{\Delta}(X, \overline{R}^{\Delta}) =& \Psi^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))] \\ &+ \psi^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)u_P(\alpha(\overline{R}^{\Delta}) - \delta),\end{aligned}$$

while our value function J^{Δ} which stops at $\overline{R}^{\Delta}$ is

$$J^{\Delta}(X, \overline{R}^{\Delta}) = \Psi^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)\nu + \psi^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)u_P(\alpha(\overline{R}^{\Delta}) - \delta).$$

By moving utility at $\tau_+(\overline{R}^{\Delta})$ to ν, we move from $\widehat{J}^{\Delta}(X, M)$ to $J^{\Delta}(X, M)$. The derivative of $\widehat{J}^{\Delta}(X, \overline{R}^{\Delta})$ with respect to X is

$$\begin{aligned}
&\Psi_X^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))] \\
&+ \Psi^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)\frac{d}{dX}\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))] \\
&+ \psi_X^{\Delta}(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X)u_P(\alpha(\overline{R}^{\Delta}) - \delta).
\end{aligned}$$

We know this is negative because $\widehat{J}^{\Delta}(X, M)$ is decreasing in X. For small Δ, we know that $\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))] \approx \widehat{J}(\overline{R}^{\Delta}, \overline{R}^{\Delta})$ and so

$$\frac{d}{dX}\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))] \approx 0.$$

Because $\Psi_X(\overline{R}^{\Delta}, S(\overline{R}^{\Delta}), X) > 0$, as we decrease $\mathbb{E}[\widehat{J}(X_{\tau_+(\overline{R}^{\Delta})}, X_{\tau_+(\overline{R}^{\Delta})})|X, \tau_+(\overline{R}^{\Delta}) < \tau_-(S(\overline{R}^{\Delta}))]$ towards ν we get that the derivative with respect to X becomes even more negative. Therefore, we have $\frac{d}{dX}J^{\Delta}(X, \overline{R}^{\Delta}) < 0$ at $X = \overline{R}^{\Delta}$. Using this observation along with continuity of P's payoffs with respect to $\overline{R}^{\Delta}$ and $J^{\Delta}(\overline{R}^{\Delta}, \overline{R}^{\Delta})$, we conclude that $J^{\Delta}(X, \overline{R}^{\Delta}) > \nu$ for all $X < \overline{R}^{\Delta}$. □

We can now formally state Theorem 3, which includes both Proposition 5 and 6 as special cases.

Theorem 3. *Let P's outside option be $\nu \geq 0$ and $J_\nu^*(X_0; W)$ be the value in the optimal mechanism which delivers W utility to A when starting at X_0. Suppose that P's continuation value is weakly above ν for all histories of this mechanism. There exists a sequence of subgame-perfect equilibrium in the discrete time alternating offers game which delivers equilibrium value J^{Δ}, W^{Δ} to P and A, respectively, such that $\lim_{\Delta\to 0} W^{\Delta} = W$ and $\lim_{\Delta\to 0} J^{\Delta} = J_\nu^*(X_0, W)$.*

Proof. Given the equilibrium structure, it is clear that the equilibrium values will converge to the continuous time limit as $\Delta, \delta \to 0$; therefore we only need verify that the proposed strategies are indeed an equilibrium.

First, we argue the outside equilibrium is sub-game perfect. The player making the proposal has no incentive to change his demand because he is getting the maximal amount he can subject to the other player's IR constraint. He has no incentive to delay because he expects to only get a value equal to his outside option at $t + \Delta$. Taking the outside option following a rejection is optimal because he expects to earn at most his outside option in the next period and hence is better off taking the outside option today.

His opponent finds it optimal to take the offer because it is IR and rejection will lead to the proposing player taking the outside option immediately.

Next, we argue that our inside equilibria are sub-game perfect. We start with an arbitrary inside equilibrium to punish P. If A deviates and increases his offer at time t when called to offer β_I, he expects it to be rejected and to receive a continuation value less than $X_t + k\delta$. Therefore, A has no incentive to raise his offer when $\beta_t = \beta_I$ if $X_t + k\delta < u_A(\beta_I)$, which holds by definition of our equilibrium. A has no incentive to raise his offer at $X_t + k\delta > u_A(\beta_I)$ because he is making the maximal demand that is IR for P. For small Δ, δ, P has no incentive to accept an offer he is called to reject because $u(\beta_{t+\Delta})$ is close or equal to ν, whereas his continuation value in the inside equilibrium to punish A is bounded away from ν. Moreover, P has no incentive to reject a demand $\beta_{t+\Delta}$ because $u_P(1-\beta_{t+\Delta}) \geq \nu$ and P expects a rejection to induce A to take the outside option immediately.

Next we consider the incentive to deviate in our on-path play or in the inside equilibrium to punish A. P has no incentive to raise his offer because he expects A to reject it and make an offer at $t + \Delta$ which leaves P indifferent between taking his outside option at t or A's offer at $t + \Delta$. A has no incentive to reject an offer he is called to accept since every offer P makes satisfies DIR. We must argue that A has no incentive to accept an offer he is called to reject. For A to accept the deviating demand α_t today, it must be that

$$u_A(1 - \alpha_t) \geq e^{-r\Delta}\mathbb{E}[max\{u_A(\beta_{t+\Delta}), X_{t+\Delta}\}|X_t] \geq e^{-r\Delta}u_A(1 - u_P^{-1}(e^{r\Delta}\nu) - \delta_P).$$

As we take $\Delta \to 0$, we have $\delta_P \to 0$. Thus, for any $\epsilon > 0$, there exists Δ small enough that $u_A(1 - \alpha_t) \geq u_A(1 - u_P^{-1}(\nu)) - \epsilon$. This implies that, for small ϵ, $u_P(\alpha_t) \leq u_P\left(1 - u_A^{-1}\big(u_A(1 - u_P^{-1}(\nu)) - \epsilon\big)\right) \approx \nu$.

But, using Lemma D.12, we know that P's continuation value is bounded away from ν except when $X_t \approx \overline{R}^{\Delta}$. Near $\overline{R}^{\Delta}$, P's continuation value will be close to ν but for $\overline{R}^{\Delta} \approx \widetilde{R}$, there will be no split which is IR for both P and A. Therefore there is no profitable deviation for P.

Our arguments imply that both players have no incentive to deviate in the offers they make or the offers they accept. Finally, we note that they have no incentive to take the outside option earlier than called to because each player's IR constraints are satisfied in the optimal mechanism and therefore will be satisfied in our approximating mechanism: P's IR constraints are strictly satisfied by Lemma D.12 while A's will be strictly satisfied in the optimal mechanism approximation since we decrease P's demand by δ.

□

Appendix E Supplementary Lemmas

Lemma E.13. *If, in $(\tau, d_\tau, \alpha_\tau)$, A's continuation value is greater than the outside option after every history, then $(\tau, d_\tau, \alpha_\tau)$ satisfies DIR.*

Proof. Take any $\tau' \in \mathcal{T}$. We can write $V(\tau, d_\tau, \alpha_\tau) - V(\tau \wedge \tau', d_\tau, \alpha_\tau)$ as

$$\mathbb{E}[e^{-r\tau'}\{\mathbb{E}[e^{-r(\tau-\tau')}(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau)|h_{\tau'}] - X_{\tau'}\}|X_0]$$

which is weakly positive A's continuation value being greater than the outside option after every history means $\mathbb{E}[e^{-r(\tau-\tau')}(d_\tau(u_A(1-\alpha_\tau) - X_\tau) + X_\tau)|h_{\tau'}] \geq X_{\tau'}$. Since this holds for all τ', DIR is satisfied. □

Lemma E.14. *The argmax over S of $F(S, W)$ is unique.*

Proof. Consider the constrained optimal problem in which we treat X^m as an absorbing barrier.

$$\sup_{(\tau,\alpha_\tau)} \mathbb{E}[e^{-r\tau}(u_P(\alpha)\mathbb{1}(X_\tau > X^m) + H(X^m)\mathbb{1}(X_\tau \leq X^m)|X] \tag{11}$$
$$\text{subject to } \mathbb{E}[e^{-r\tau}(u_A(1-\alpha)\mathbb{1}(X_\tau > X^m) + X^m\mathbb{1}(X_\tau \leq X^m)|X] \geq W$$

Because this problem allows for richer mechanisms than just a static threshold as in $F(S, W)$, its value will yield an upper-bound on $F(S, W)$. We will argue that they are equal by showing the solution to 11 takes the form of a static lower threshold and that this threshold is uniquely determined.

We start by transforming 11 into a Lagrangian. There exists multiplier $\eta < 0$ such that τ is a solution

$$U(X) = \sup_{(\tau,\alpha_\tau)} \mathbb{E}[e^{-r\tau}(u_P(\alpha) - \eta u_A(1-\alpha))\mathbb{1}(X_\tau > X^m) + (H(X^m) - \eta X^m)\mathbb{1}(X_\tau \leq X^m)|X].$$

Because the optimal choice of α_τ is always $\alpha^* = \max_{\alpha\in[0,1]} u_P(\alpha) - \eta u_A(1-\alpha)$, we can rewrite U as

$$U(X) = \sup_{\tau} \mathbb{E}[e^{-r\tau}\{u_P(\alpha^*) - \eta u_A(1-\alpha^*))\mathbb{1}(X_\tau > X^m) + (H(X^m) - \eta X^m)\mathbb{1}(X_\tau \leq X^m)\}|X]. \tag{12}$$

By Balzer and Janßen (2002), any solution to 11 will be a solution to 12. Therefore we are done if we can show there is a unique solution to 12.

By similar arguments as in Lemma A.2, $U(X') \geq u_P(\alpha^*) - \eta u_A(1-\alpha^*))$ for all $X' < X^m$ and the solution to 12 takes the form $\tau = inf\{t : X_t \notin (S^*, R)\}$ for some S, R.

We start by arguing that $R = X^m$. If this were not the case, then it would be better to stop immediately (i.e., $S = X^m$) since $u_P(\alpha^*) - \eta u_A(1-\alpha^*)$ is independent of X. Next, we argue that τ is unique. Suppose there was another optimal rule $\tau' = inf\{t : X_t \notin (S', X^m)\}$ with $S' < S^*$ (the arguments will be similar if $S^* < S'$). Continuation value, under either rule, at any X'' will equal $U(X'')$. Because it is not better to stop immediately at S^* when using τ', standard dynamic programming arguments imply that the continuation value at $X' \in (S', S^*)$ is

$$\begin{aligned} U(X') &= \mathbb{E}[e^{-r(\tau_-(S')\wedge\tau_+(S^*))}U(X_{\tau_-(S')\wedge\tau_+(S^*)})|X'] \\ &= \mathbb{E}[e^{-r(\tau_-(S')\wedge\tau_+(S^*))}\big(u_P(\alpha^*) - \eta u_A(1-\alpha^*)\big)|X'] \\ &< u_P(\alpha^*) - \eta u_A(1-\alpha^*), \end{aligned}$$

which contradicts $U(X') \geq u_P(\alpha^*) - \eta u_A(1-\alpha^*)$.

□