Brook no compromise: How to negotiate a united front

Elaine Yao*

March 28, 2024

Abstract

Negotiating factional conflict is crucial to successful policymaking. In these conflicts, actors sometimes employ hardball tactics to strategically rule out outcomes they dislike. Using a dynamic bargaining model, I explore how the threat and usage of these tactics impact coordination between actors with conflicting interests. In the model, two players who prefer different reforms must jointly agree on one to overturn a mutually unfavorable status quo. Neither knows whether the opponent prefers the status quo over their less-preferred outcome. Players willing to compromise rationally delay agreement, balancing the incentive to preempt the opponent against the benefit of waiting to gather better information. Delay is prolonged when players cannot easily glean one another's willingness to compromise. I show that when private willingness to compromise is likely to be revealed, players delay longer. Thus, higher-leak environments are beneficial to welfare, as the additional delay incentivized by leaks deters mistakes of preemption.

*I thank Gérman Gieczewski and Matías Iaryczower for guidance throughout this project. I am grateful for the discussions and feedback from seminar audiences at Princeton, Virginia Tech, and participants at meetings of the Southern and Midwest Political Science Associations.

1 Introduction

Playing hardball is a dangerous, yet ubiquitous political tactic. Sparring rebel factions kill one another's leaders, either allowing one faction to dominate the movement or dooming their collective cause. Factions within and across political parties kill one another's bills, either guaranteeing the passage of more favorable legislation or prolonging legislative gridlock. Moral hazard hawks and liquidity doves veto one another's bailout packages, either averting calamitous financial mismanagement or allowing financial crises to propagate out of control. When do organizations make such irrevocable commitments, and what holds them back? When do these tactics help one side get its way, and when do they backfire to the detriment of both sides?

When hardball tactics do backfire, they can cause negotiations to fail in ways that seem accidental and even preventable. Within the U.S. House of Representatives, the ultraconservative "Freedom Caucus" has developed a reputation for employing hardball tactics against the moderate wing of its own party. In 2017, Republican House leadership sought to overturn the Affordable Care Act, a policy long reviled by both moderate and conservative wings of the party. Confident that they would be able to win the support of the Freedom Caucus, moderate leaders brought a partial repeal bill to the floor. However, Freedom Caucus members who had called for a more radical rollback of the ACA withheld the votes necessary to pass the bill they had decried as "Obamacare Lite" (Bade, Dawsey and Haberkorn 2017). Caught unexpectedly short of votes, leaders were forced to pull the bill from the floor, allowing the ACA to remain in place.

The essential quality of hardball tactics is to render the opponent's preferred outcome unobtainable. When one faction brings their preferred legislation to the floor, they force their opponent into a choice between accepting that legislation or rejecting it. When the European Central Bank held the IMF hostage over the issue of Greek debt relief, it essentially forced the IMF to choose between signing onto what the IMF saw as a suboptimal rescue plan, and no plan at all (James 2024). When snipers opened fire onto a crowd gathered to welcome Juan Perón's return to Argentina from exile, they sought to intimidate left-wing Peronists into yielding to right-wing leadership of the movement, marking the beginning of Argentina's Dirty War (Horowicz 2009).

Whether due to electoral promises, financial calculations, or ideological convictions, actors can and do interrupt the usual back-and-forth of negotiations to take one option off the negotiating table. Departing from the conventional bargaining framework, I study how the availability of these tactics shape players' beliefs about one another's willingness to compromise as well the likelihood of negotiation failure. I develop a dynamic game with two players who must jointly agree to one of two possible reforms in order to overturn a status quo which both dislike. Each player prefers a different reform, but neither initially knows the other's willingness to compromise: that is, what their opponent would choose if forced to choose between their less-preferred policy and the status quo. Over time, players receive opportunities to play hardball and permanently eliminate one possible reform. If they do so, they leave their opponent with a choice between a potentially sub-optimal policy and the status quo.

Akin to the textbook Battle of the Sexes, this game hinges on the tension between conflict and coordination. In the unique equilibrium, "hard" types that would never support their opponent's policy play hardball as soon as they can, forcing their opponents to choose between their own favorite policy and the status quo. By contrast, "soft" types, who are willing to compromise on their opponent's outcome to avoid the status quo, stall strategically to learn about their opponent's willingness to compromise. They cannot stall indefinitely, however, as doing so risks being preempted by their opponent and forced into a suboptimal compromise. However, too-short delay could permanently alienate an obstinate opponent and result in an avoidably poor outcome, analogous to when Republican leaders preemptively pushed forward a repeal bill that Freedom Caucus members would not support.

Such mistakes of preemption are possible because groups cannot simply play hardball whenever they please. In reality, frictions in intra-group communication and discipline, the cost of acquiring and mobilizing resources, institutional procedures, idiosyncratic political opportunities, and other constraints mean that groups cannot instantaneously mobilize troops, respond to financial crises, or put policies to a vote. Reflecting this, players in the model can only play hardball when they receive an opportunity to do so. These opportunities arrive stochastically, meaning that even if a group is ready and willing to play hardball from the start, it may still be preempted by its opponent if it does not receive a commitment opportunity in time. Soft types who choose to play hardball before discovering their opponents' true type knowingly take on the risk of triggering an avoidable stalemate. Since avoidable miscoordination is most likely when delay is short, factors which prompt soft types to act more preemptively – such as prior beliefs about opponents, policy priorities, and the frequency of hardball opportunities – also increase the likelihood of avoidable miscoordination.

Not all soft types are created equal. Even if both players are soft, they may have different costs of compromise due to policy priorities or electoral pressures, or different abilities to play hardball due to group size or group discipline. In an asymmetric setting of the model, I demonstrate that when model parameters can vary between types *and* players, the shape of equilibrium remains largely unchanged. In fact, some strategic considerations come into clearer focus: When a soft type of one player is less inclined to delay hardball, a soft type of its opponent delays even longer. The reason is that once a player races to commit irrespective of its type, its opponent can no longer make inferences about its willingness to compromise from type-specific commitment behavior. The opponent's slower learning translates into longer delay, in turn deterring mistakes of preemption that cause avoidable miscoordination. Therefore, the symmetric setting presents a "worst-case scenario" for the likelihood of failed negotiations, as any amount of symmetry induces longer delay.

Delay is costly for another reason: the risk of private and potentially damaging information coming to light. When assumed-to-be-private conversations, memos, and drafts are unexpectedly leaked, they can directly reveal or pressure parties into revealing the strength or weakness of their position. This can expose previously secret willingness to compromise, and sometimes change the direction of negotiations. In negotiations between the IMF and European creditors over the issue of debt relief, the leak of an internal phone call wherein officials threatened to walk away from talks if European officials failed to make concessions on debt relief prompted European officials to backpedal in order to not lose IMF participation and further destabilize the situation (Dendrinou and Steinhauser 2016; Walker and Stamouli 2016). Do the occurrence or mere threat of leaks motivate actors to act preemptively, hoping aggression will compensate for their informational disadvantage? Or do actors slow down to avoid making mistakes against potentially better informed opponents?

The model yields a clear answer: when a player's willingness to compromise is exposed, they delay longer. The logic mirrors that of the asymmetric setting: once a leak occurs, the opponent of the leaked player has an incentive to play hardball irrespective of their type. Prior to a leak occurring, a soft opponent would have held back, but afterwards, they no longer have a reason to wait, and assume a posture identical to a hard opponent. While the leaked group's chances of implementing its preferred policy fall, the chances of a failed negotiation also fall. In fact, simply increasing the average rate of leaks occurring causes soft players to delay more, even without a leak ever actually occurring. Substantively, these results suggest that in higher-leak environments, players exhibit cautionary behavior that reduces the incidence of avoidable miscoordination, but also lengthens the duration of negotiation.

The paper is organized as follows. Section 2 discusses related theoretical and empirical literature. Section 3 describes the model. Section 4 characterizes equilibrium in the symmetric setting. Section 5 analyzes welfare in the symmetric setting. Section 6 generalizes results to the asymmetric setting. Section 7 concludes. All proofs can be found in the Appendix.

2 Literature

Factional contestation has historically been studied through the lens of delegation and power-balancing. For instance, scholars of American congressional politics have long debated the extent to which leaders are able to use procedural controls and distributions of rewards to control rank-and-file members (Cox and McCubbins 2005; Kiewiet and McCubbins 1991; Krehbiel 2010). In nondemocratic and democratic transition contexts, contestation emerges from power balancing between dictators, elites, and civil society actors (De Mesquita et al. 2005; Meng, Paine and Powell 2023; O'Donnell and Schmitter 1986; Svolik 2009). In both cases, factional conflict has long taken a backseat to such questions about the particular institutional structures that enable certain parties to hold onto power and advance their interests.

A growing body of work, particularly within American politics, has begun to shift the lens onto internecine conflict itself. For instance, Rubin (2017) studies the history of intra-party organizations in order to elucidate their role in solving coordination problems and gaining leverage over non-allied party leaders, and Green (2019) empirically studies the unconventional threat-making tactics employed by the House Freedom Caucus. Formal theorists have begun to study these questions as well, with recent papers by Izzo (2020) and Invernizzi (2023) describing the electoral incentives that shape factional contention. Unraveling the dynamics of factional contestation complements the longstanding literature on how democratic and nondemocratic actors alike bargain for political survival and push through their desired outcomes. As my model makes clear, different notions of "power" can have vastly different effects on negotiation tactics. Perceived differences in strength may be reflected in players' priors about one another's types, the frequency with which players receive commitment opportunities, or the frequency with which factions uncover compromising information about their opponents. As my model shows, these do not impact players' behavior in identical ways.

Bargaining is the most widely employed formal theoretic approach to analyzing negotiations. For instance, a long literature on crisis bargaining in international relations, originating with Fearon (1994), employs reputational bargaining models in the vein of Abreu and Gul (2000) and Fudenberg and Tirole (1986). In these models, players begin with high distributional demands, from which 'rational" players concede at a constant rate over time, while "irrational" behavioral types hold out forever for their high demand.

The war of attrition dynamics inherent in these games are reversed by equilibrium learning structure of my game. This distinction arises from my choice to model actions as commitments from a low status quo, rather than concessions from a high demand. In traditional reputational bargaining models, players progressively become sure that, in the absence of concessions, their opponents are likely to be irrational types. In my model, players develop increasing certainty over time that, without any commitments, their opponents are more likely to be "soft" compromising types.[1] Preemptively committing to policy in order to bluff as a hard type can be extremely costly for a group that is not prepared to accept the status quo as an outcome of their action. This observation underlines the importance of the assumption to model hardball tactics as permanent commitments to taking one option off the table entirely. Miscoordination is the risk that players inevitably take on when they employ such tactics. By contrast, in reputational bargaining models, miscoordination is never "locked in" – in theory, the door to agreement always remains open.

Unlike bargaining games that strictly structure the sequence of play, I give both players the choice of using hardball tactics in continuous time, although the rate at which players receive opportunities to use these tactics may differ. This allows for dynamic predictions about delay and preemption to arise that are absent from other bargaining games that seek to analyze take-it-or-leave-it offers, including veto bargaining games in American politics and ultimatum games in crisis bargaining (Cameron and McCarty 2004; Fey and Kenkel 2021). Trade-offs between preemption and caution exist outside of this literature, but are not underpinned by a coordination motive. Weeds (2002) models strategic delay undertaken by firms seeking to make a one-time, irreversible R&D investment who face both a preemptive incentive (the winner-takes-all nature of the patent system) and a caution incentive (uncertainty over the profitability of the investment). In Weeds' model, firms are similarly "stuck" once they make an investment, but the outcome of the investment is unknown due to economic and technological uncertainty. The motive for delay is thus uncertainty over the profitability of the investment, rather than uncertainty about the cooperativeness of the rival firm.

Some features of the model reflect recent approaches in economic theory. By modeling the arrival of commitment opportunities as a Poisson process, I follow Kamada and Sugaya

[1] It is useful to view the game I develop as a dynamic extension of the Battle of the Sexes, and reputational bargaining/war of attrition as a dynamic extension of Chicken. While the static versions of these games are equivalent up to swapping the names of actions, they generate different incentives and patterns of learning in the dynamic versions.

(2020) and related papers in the revision games literature, e.g. Calcagno et al. (2014); Kamada and Kandori (2020).[2] This approach bears similarities to that of Ambrus and Lu (2015), who develop a continuous-time coalition bargaining game in which players get Poisson-distributed opportunities to make proposals to one another. These frictions allow for comparative statics that answer key questions, such as the effect of a higher-leak environment on the length of delay and incidence of miscoordination.

3 Model

There are two infinitely lived players, a and b, which I refer to as groups. Time $t \in [0, \infty)$ is continuous. There is a status quo (SQ) policy in place at the start of the game, as well as two alternatives, A and B. It is public knowledge that group a's preferred policy is A, and group b's preferred policy is B. Policy can only be changed once in the game and requires the consent of both groups. Each group is either a "soft" or "hard" type. Hard types prefer SQ to their opponent's preferred alternative, and soft types prefer their opponent's preferred alternative to SQ. Let $u_i^\theta(X)$ denote the utility of a group $i = a, b$ of type $\theta = s, h$ for some alternative $X \in \{A, B, SQ\}$. Then, for a hard type of group a,

$$u_a^h(A) > u_a^h(SQ) > u_a^h(B)$$

For a soft type of group a,

$$u_a^s(A) > u_a^s(B) > u_a^s(SQ)$$

Hard and soft types of group b satisfy analogous properties. As I will show that hard types act largely mechanically, I will omit the s superscript on the utility functions of soft types in equilibrium analysis.

Groups stochastically receive *commitment opportunities* at which time they can choose to render one policy unobtainable. These rate of these arrivals is type-specific. A group of type θ receives commitment opportunities at rate Poisson(μ_θ). Arrivals of opportunities are private information, and remain private if a group chooses to pass. If a group acts on an opportunity by committing to a policy, its opponent must immediately accept or reject. If the opponent accepts the policy, the policy is implemented, and if it refuses, the status quo remains in place. Once a group acts on an opportunity, its decision is permanent and irrevocable. This is what I refer to as hardball tactics. After the opponent's decision, the game ends and players receive

[2] As Kamada and Sugaya observe, this approach is analogous to the technique commonly used in macroeconomics to model uncertainty over future opportunities to change prices originated by Calvo (1983).

their infinite-horizon utility from the final policy that is implemented. There is no discounting.

A group's type may be publicly revealed or "leaked" at any point. The stochastic arrival rate of information leaks is type-dependent: the probability that a group of type θ is leaked is Poisson(λ_θ). I assume that $\mu_h + \lambda_h \geq \mu_s + \lambda_s$, that is, the combined average rates at which hard types are leaked or can commit exceeds that at which soft types are leaked or can commit.[3] At the start of the game, group a holds a prior belief $p_b \in [0, 1]$ that their opponent b is a hard type. Analogously, p_a is the prior about opponent a held by group b. In the absence of leaks, groups update their beliefs endogenously based on opponents' actions. The solution concept is Perfect Bayesian Equilibrium.

3.1 Discussion of model assumptions

In the initial presentation of the model, I assume symmetry across groups. That is, for $\theta \in \{s, h\}$, I assume $p_a = p_b \equiv p$, and $u_a^\theta(A) = u_b^\theta(B)$, $u_a^\theta(B) = u_b^\theta(A)$, and $u_a^\theta(SQ) = u_b^\theta(SQ)$. I also assume throughout that groups do not discount the future. This allows me to isolate groups' incentives to preempt undiluted by the additional impatience generated by discounting. Therefore, if policy changes at any point, the new policy fully dominates infinite-horizon utility. In the absence of discounting, it is technically necessary to impose a tie-breaking assumption for situations where players are indifferent between committing at time t and committing at any time after t. I will assume throughout that if players are indifferent between committing at t and any time after t, they will commit at t. If players discount the future by even an infinitesimal amount, they would prefer to commit at t.

The exposition of the model suggests that if neither group's preferred reform is passed, the original status quo will be instated. However, the logic of the model remains unchanged if we take the status quo to represent a different outcome that is instated if an offer is made and refused. Indeed, it is the status quo that is instated *after an unsuccessful offer* that factors into players' strategic calculus. Therefore, the model remains identical if SQ is replaced with a different option instated only if negotiation fails. In some instances, it is plausible to assume that this is identical to the status quo that players begin with (e.g. when Obamacare remained in place before and after the unsuccessfully attempted repeal). In others, it is more plausible to assume this is something else (e.g. civil war, which is arguably more detrimental than a previously existing alliance).

[3]The equilibrium characterization in threshold strategies holds when this condition holds with equality, but not if $\mu_s + \lambda_s > \mu_h + \lambda_h$. The utility of this assumption becomes clear after the equilibrium characterization, and is further addressed in the Discussion.

The model imposes weak assumptions on preferences. It does not require that players have single-peaked preferences, although it nests this possibility. As an example, suppose players have single-peaked preferences and $SQ < A < B$, that is, B is the most ideologically extreme far-right reform and A is a more moderate right-leaning reform, but the exact policy locations of reforms are uncertain. In the absence of any non-ideological dimensions of policy that players care about, such as quality or public approval, then b must be a soft type. Whether a is a hard or soft type depends on whether A is closer to B or to SQ. This case where one player's type is known and the other's is unknown is captured by Proposition 2, where the soft group's type is revealed at $t = 0$. However, the ideological location of policies is rarely the sole factor relevant to groups' decisions. Policy quality aside, intertemporal and reputational concerns may also induce players to act against what single-peaked preferences would apparently dictate. For instance, if b stonewalls today's policy, they may increase p_b and bolster their position in a future higher-stakes negotiation. While I do not model these idiosyncratic intertemporal concerns explicitly, I emphasize that their presence in negotiations means that whether a group is "hard" or "soft" on any given issue is not fully determined by policy preferences alone.

4 Equilibrium

I begin by describing the full information benchmark where types are public. In this benchmark, groups' optimal strategies never involve delay. There is no uncertainty to be resolved by delay, so soft types make offers as soon as they receive an opportunity. Since groups' types are publicly known, there is also no efficiency loss from avoidable miscoordination.

Remark 1 (Full information benchmark)**.** *If both groups are hard types, the status quo is never overturned. If both groups are soft types, the first group to receive a commitment opportunity determines the final policy. If groups are different types, the final policy is the one preferred by the hard type.*

In the full version of the model with incomplete information there is both delay and efficiency loss. Hard types still have a dominant strategy to propose their preferred alternative as soon as they receive an opportunity. Since their optimal strategy is independent of their beliefs or their opponents' beliefs, they act mechanically. Soft types that know their opponent's type also have a dominant strategy to propose a policy as soon as possible. Soft types that have full information act identically to the full information benchmark: they commit to their preferred alternative if the opponent is a soft type, and to their opponent's preferred alternative if the opponent is a hard type.

The more interesting problem, which motivates the first proposition, is faced by soft types who do not know their opponents' type. To begin the analysis, suppose that neither player's type has been leaked up to some time t. I refer to this as a *relevant history.* At a relevant history, a soft type ("the decisionmaker"; she) knows that a hard type of her opponent has been trying to commit from the start. As more time passes without her opponent playing hardball, the less the decisionmaker believes that her opponent is a hard type. A soft opponent, meanwhile, is making the same calculation as the decisionmaker, becoming more and more convinced over time that the decisionmaker is a soft type and therefore safe to preempt. At a certain time, the decisionmaker's cost of being preempted by a potentially soft opponent will outweigh the benefit of giving a potentially hard opponent sufficient time to act first. This result is formalized below:

Proposition 1 (Commitment delay for soft types)**.** *Define a relevant history as one where neither group has committed to a policy and no group's type has been leaked. Consider the continuation game at a relevant history. There exists a unique threshold time T^* such that if both groups' types remain private information, a soft group that receives a opportunity at time t will commit to its preferred alternative if and only if $t > T^*$.*

Proposition 1 describes the first time at which a soft type is willing to play hardball. This time corresponds to a particular posterior belief which equalizes the expected value of delay and expected value of making a permanent commitment. Mathematically, this threshold belief is at the core of all equilibrium expressions and encodes many of the factors which tilt players' behavior towards preemption or delay. For a soft type of player a, this threshold belief is:

$$\frac{\mathbb{P}\left(b \text{ is a hard type}\right)}{\mathbb{P}\left(b \text{ is a soft type}\right)} = \frac{1}{2}\frac{(\lambda_h + \mu_h) + \mu_s}{(\lambda_h + \mu_h)}\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \quad \equiv P_T \tag{1}$$

In P_T, $\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)}$ describes the strength of a's preference for policy A over B, relative to B over SQ. I call this *relative desirability.* When relative desirability is high, B is less substitutable for A, and as a result P_T is higher. This functions like a marginal rate of substitution: when relative desirability $= 1$, A is as preferable to B as B is preferable to SQ. As relative desirability increases, A becomes much better than B, and as it decreases, A is only slightly better than B. As relative desirability $\to \infty$, a soft type becomes closer to a hard type. Accordingly, if relative desirability increases, P_T also increases, becoming an *easier* threshold to cross (as beliefs are decreasing from the prior) and tilting behavior towards preemption. Note as well that P_T depends on the sum of λ_h and μ_h. From the perspective of a soft type of a, regardless of whether its opponent is leaked as a hard type or commits to policy B, the effect is the same: Policy A is impossible to get.

Visualization of Proposition 2

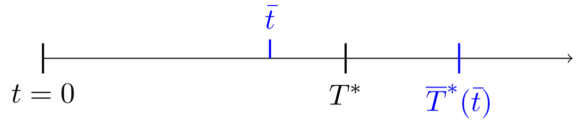

Figure 1. T^* is when the soft group becomes willing to act if both groups' types are unknown. If the soft group's type is revealed at $\bar{t}$, it begins committing starting at $\overline{T}^*$, the value of which is dependent on $\bar{t}$.

Using the threshold belief P_T, we can complete equilibrium characterization. Proposition 1 is only a partial characterization, as it assumes that *neither group* has had their type revealed. When a decisionmaker is revealed to be a soft type, she behaves differently in the face of a fully informed opponent:

Proposition 2 (Delay with asymmetric information)**.** *Suppose a soft group's type is revealed at time $\bar{t} < T^*$. Then, there exists a unique $\overline{T}^*(\bar{t}) > T^*$ such that a soft group will commit to its preferred alternative iff $t > \overline{T}^*(\bar{t})$. Moreover, $\overline{T}^*(\bar{t})$ is decreasing in $\bar{t}$, with $\overline{T}^*(T^*) = T^*$.*

Furthermore, T^* and $\overline{T}^*$ are given by

$$T^* = \max\left\{\frac{1}{\lambda_s - (\lambda_h + \mu_h)} \ln\left(\frac{1-p}{p} P_T\right),\ 0\right\} \tag{2}$$

$$\overline{T}^* = \max\left\{\frac{1}{\lambda_s + \mu_s - (\lambda_h + \mu_h)} \left[\ln\left(\frac{1-p}{p} P_T\right) +\ \mu_s \bar{t}\right],\ 0\right\} \tag{3}$$

This result suggests that we should observe *more* delay when groups have asymmetric information. After group a is revealed to be a soft type, the opponent b wants to commit as soon as possible regardless of b's type. From a's perspective, this means that *inaction is no longer informative about b's type.* This shuts down one channel through which a updates beliefs. Since a is now learning more slowly, it takes more time for posteriors to converge to P_T. Importantly, P_T itself remains unchanged – only the speed of convergence changes. Furthermore, $T^* = \overline{T}^*(T^*)$, that is, if a soft type is leaked exactly at the threshold time T^*, there is no delay as their beliefs have already reached the threshold.

What factors impact group behavior? We have already observed that high relative desirability induced preemption by making soft types less willing to compromise. Similarly, decreasing a group's prior belief that their opponent is hard also incentivizes preemption. Intuitively, this reduces the distance between prior beliefs and P_T. I summarize these two results as follows:

Corollary 1 (Conditions for no delay)**.** *A soft type of group a acts without delay when one or more of the following are true:*

1. *The relative desirability of A compared to B is high, i.e.* $\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)}$ *is high enough.*
2. *Group a has a strong prior belief that b is a soft type, i.e.* $\frac{1-p}{p}$ *is high enough.*

Changing the rates of leaks and commitment opportunities is technically and substantively subtler. These comparative statics speak to questions about the information and technological environment that players inhabit: What happens when players expect their groups' internal discussions to be leaked more frequently owing to an increasingly high-frequency news cycle and demand for journalistic investigations? What happens when groups' military and communications technology improves to the extent that they can more rapidly agree amongst themselves to deploy force?

The answers are not always mathematically straightforward. Changing rates may affect both threshold beliefs *and* the rate of convergence. When these effects move in opposite directions – for instance, threshold beliefs become easier to reach, but the rate of convergence is slower – comparative statics can sometimes be non-monotonic. To begin the analysis, note that increasing $\lambda_h + \mu_h$, the effective rate at which hard types screen out, leads to less delay by soft types. With this straightforward result in hand, I hold constant λ_h and μ_h in order to understand the effects of changing rates that apply to soft types.

Increasing the rate at which soft types are leaked, λ_s, slows the rate of learning, leading to more delay. Although it seems counterintuitive that increasing the incidence of leaks makes it more difficult to learn, the key is that in the absence of any actual leaks, soft types learn on the basis of the difference $\lambda_h - \lambda_s$, in other words, how frequently they should expect a hard type of their opponent to be leaked relative to a soft type. First suppose this difference is large. Then, observing no leaks for a long duration of time is strong evidence that the other player is soft. If this difference is small, then observing no leaks is not very informative, and a soft type must wait much more time to become confident the opponent is a soft type. In other words, when λ_s increases, learning is more difficult, extending delay.

The rate at which soft types receive commitment opportunities, μ_s, has a more ambiguous impact on delay, since it can change the rate of learning as well as the beliefs threshold. On one hand, increasing μ_s increases P_T. In the case where no group has been leaked, this is the *only* place in which μ_s appears in equilibrium delay. Thus, T^* – the soft type's threshold delay in the no-leak equilibrium – decreases in μ_s. However, μ_s has another impact in the case where a soft type has been leaked. To see the result intuitively, recall that after the leak, both types of opponent pool on actions. Thus, the leaked soft type must now make inferences based on the difference $\mu_h - \mu_s$. Similarly to the analysis of λ_s, increasing μ_s makes this difference smaller, and impedes learning.

Whether $\overline{T}^*$ is increasing or decreasing in μ_s depends on which of these two effects dominates: the "lower bar" effect on P_T which incentivizes preemption, or the "slower learning" effect which incentivizes delay. If the soft type is leaked early, slower learning takes place over nearly the entire duration of the game, tending to dominate the change to P_T. By contrast, if the soft type is leaked late in the game, the majority of learning about the opponent's type is already done before the leak. Slowing the speed of remaining learning has an inconsequential effect relative to the bump to P_T, which is sufficient to reduce delay. In the asymmetric case, therefore, the comparative static on μ_s is nonmonotonic.

I summarize these observations below:

Proposition 3 (Comparative statics for delay and rate parameters)**.**

1. *T^* and $\overline{T}^*$ are increasing in λ_s, the rate at which soft types are exogenously leaked.*

2. *T^* is decreasing in μ_s, the rate at which soft types receive commitment opportunities.*

3. *$\overline{T}^*$ is increasing in μ_s when $\bar{t} < \max\left\{T^* - \frac{\lambda_s+\mu_s-(\lambda_h+\mu_h)}{(\mu_s+(\lambda_h+\mu_h))(\lambda_s-(\lambda_h+\mu_h))}, 0\right\}$ and decreasing in μ_s otherwise.*

4. *T^* and $\overline{T}^*$ are decreasing in $\lambda_h + \mu_h$, the rate at which hard types screen out.*

This analysis makes clear how changing rates can have counterintuitive consequences. Especially in the absence of *observed* arrivals of commitments or leaks, simply knowing the *average* rate of arrivals can give players the information they need to deduce their opponent's willingness to compromise over time. Observe also that leaks have both direct and indirect effects on behavior. Besides putting the leaked player at a direct disadvantage, they also change the inferences that the leaked player can make on the basis of its opponent's behavior. In the asymmetric setting, the caution that this incentivizes for the leaked player can also give its opponent enough "slack" to delay as well.

I conclude the equilibrium characterization with a comment regarding uniqueness, which motivates results that I derive in the asymmetric setting. T^* and $\overline{T}^*$ are unique in the space of threshold strategies, that is, assuming that an uninformed soft type of the opponent is committing after a single threshold, then the best response is also threshold. Figure 2 plots each

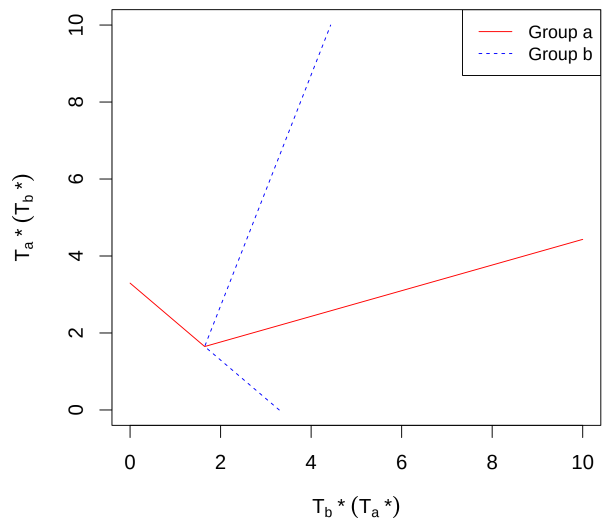

Figure 2. Example of best response correspondences in a symmetric setting.
Parameters: $\lambda_s = \frac{1}{3}, \mu_s = \frac{1}{3}, \lambda_h + \mu_h = 1, \frac{1-p}{p} = \frac{1}{4}, \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{u_b(B)-u_b(A)}{u_b(A)-u_b(SQ)} = \frac{1}{2}$

player's optimal threshold taking the opponent's threshold time as given. Each player's best response, viewed independently, takes a "V" form. In the symmetric setting, the equilibrium corresponds to a single point of intersection on the 45-degree line where $T_a^* = T_b^*$.

Consider the shape of the best responses. The point of the "V" represents the minimum possible amount of delay given a particular best response function. The symmetric case, which I have analyzed thus far, corresponds to the special case where best responses kink in the same place and therefore intersect there, at the point of both "V"s. Moving along either player's best response function away from the kink amounts to increasing delay for *both* players. I discuss this further in Section 6, where I also consider comparative statics that vary a parameter for a particular type of group a without varying it for the same type of group b. In particular, increasing the rate at which one player is leaked only shifts one best response while holding the other constant, resulting in both players delaying longer than in the symmetric case.

5 Welfare

I have described players' propensity to employ or refrain from hardball tactics, and how soft players adjust their delay in response to exogenous factors. I now turn to how equilibrium behavior maps onto players' welfare. Retaining the symmetric setting analyzed thus far, I begin by developing the connection between equilibrium strategies and welfare. Consider a soft type of player a's infinite horizon expected utility, which corresponds to how well-off they can expect to be at the start of the game:

$$(1-p)\left[\frac{u_a(A)+u_a(B)}{2}\right]+(p)\Big[\mathbb{P}\,(\text{avoidable miscoordination})\,u_a(SQ)+ \\ \Big(1-\mathbb{P}\,(\text{avoidable miscoordination})\Big)u_a(B)\Big] \quad (4)$$

Avoidable miscoordination embeds soft types' equilibrium choice of delay. Equilibrium considerations do not appear in the first term, because symmetric soft types delay the exact same amount of time and receive commitment opportunities at the same rate – hence, the resulting policy (A or B) is a coin toss. However, equilibrium considerations determine the likelihood that a soft type commits to hardball tactics before its opponent is revealed as a hard type, triggering avoidable miscoordination. As such, avoidable miscoordination is the sole strategic component of *ex ante* welfare.

Lemma 1 (Commitment delay and avoidable miscoordination)**.** *Suppose that uninformed soft types pass on all opportunities to commit until an arbitrary time T, after which they commit to their most-preferred alternative as soon as they receive a commitment opportunity. Then, the probability of avoidable miscoordination is*

$$e^{(-\mu_h-\lambda_h)T}\frac{\mu_s}{\mu_s+\mu_h+\lambda_h} \quad (5)$$

which is decreasing in T.

Lemma 1 observes that avoidable miscoordination is mechanically decreasing in delay. It is not an equilibrium statement. The equilibrium probability of avoidable miscoordination involves *ex ante* expected equilibrium delay. Computing this becomes complex, as it involves both the likelihood of no leak occurring prior to T^*, and the likelihood at every instant $\bar{t}$ that a leak occurs that will trigger the revised threshold, $\overline{T}^*(\bar{t})$. I relegate the statement to the Appendix (Equation A19).

Priors and relative desirability have mechanical as well as strategic effects on welfare. While increasing relative desirability shortens delay (making avoidable miscoordination more likely), it also mechanically boosts utility in the cases where there is no avoidable miscoordination. Similarly, increasing p means that avoidable miscoordination is less likely if the opponent is truly a hard type (increasing welfare), but makes it also more likely *a priori* that the opponent could be a hard type (decreasing welfare). Figure 3 illustrates these dynamics. Note that even in the region where the probability of avoidable miscoordination is flat (due to the fact that

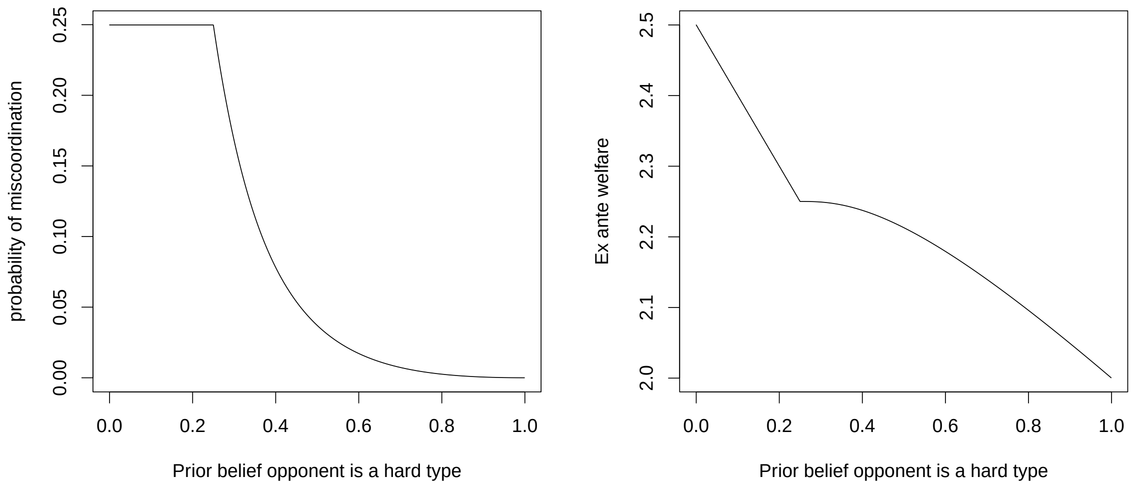

(a) Miscoordination

(b) Welfare

Figure 3. Effects of changing priors on miscoordination and welfare. Left of the kink, soft players do not delay at all. Right of the kink, soft players choose positive delay.
Parameter values: $\lambda_s = \frac{1}{3}, \mu_s = \frac{1}{3}, \lambda_h + \mu_h = 1, \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{1}{2}$

delay is 0), welfare is still falling due to the mechanical effect of the opponent being more likely to be a hard type.

By comparison, rate parameters only have strategic effects, since they only appear in *ex ante* welfare via the probability of avoidable miscoordination. Delay is increasing in λ_s, so welfare is increasing in λ_s. This comparative static is visualized in Figure A1. Stepping back from the model, the result is far from obvious. Leaks impose a cost upon soft players who delay, and yet these players are better off when this cost increases! The key intuition is that avoidable miscoordination is the only conduit between the "leakiness" of the environment and welfare. When λ_s is higher, players make fewer mistakes of preemption, and are collectively better off for it.

The effect of commitment opportunities is mediated by the likelihood of leaks. In short, higher μ_s negatively impacts welfare *as long as leaks are not too likely.* To see why, recall the effects of μ_s on delay: While T^* was decreasing in μ_s, $\overline{T}^*$ was decreasing in μ_s if a player was leaked late enough. In their *ex ante* calculation of welfare, players cannot anticipate whether or not a leak will occur early or late. They can only anticipate how they can react in any cases that a leak occurs. If μ_s is low relative to μ_h, hard types screen out discernably faster, reducing

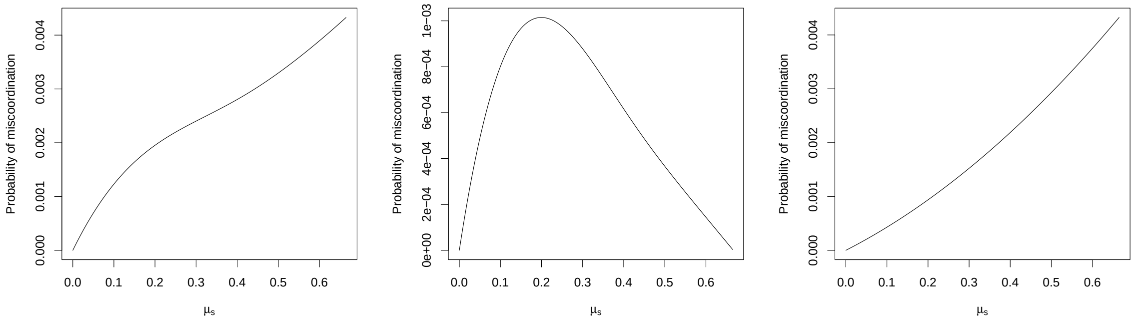

(a) Aggregate effect

(b) Effect in $\overline{T}^*$ case

(c) Effect in T^* case

Figure 4. Effects of changing μ_s on miscoordination. Left panel presents the aggregate effect, which is the probability-weighted sum of the effects in the one-sided and two-sided asymmetric information cases (panels b and c, respectively).
Parameter values: $\lambda_s = \frac{1}{3}, \lambda_h + \mu_h = 1$, $\frac{1-p}{p} = \frac{1}{4}$, $\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{1}{2}$

the incidence of miscoordination. At high values of μ_s, soft types delay a long time, since it is difficult to learn on the basis of $\mu_h - \mu_s$. *Miscoordination is most likely at intermediate values of* μ_s, as depicted in Figure 4 panel (b). When leaks are highly probable, the $\overline{T}^*$ dominates, so welfare has a "U"-shape in μ_s. When leaks are rare, the T^* case dominates, so welfare simply decreases in μ_s (Figure A2).

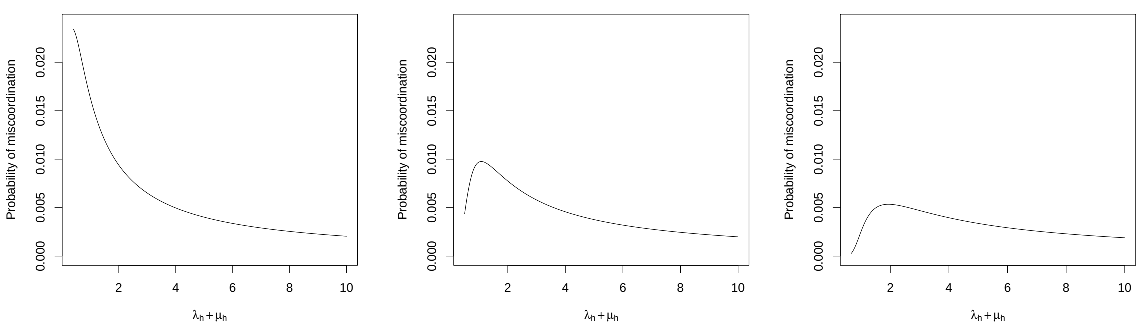

(a) $\lambda_s = \frac{1}{16}$

(b) $\lambda_s = \frac{1}{6}$

(c) $\lambda_s = \frac{1}{3}$

Figure 5. Effects of changing $\lambda_h + \mu_h$ on the probability of avoidable miscoordination. As I progressively increase λ_s across the panels, non-monotonicity in the comparative static becomes more pronounced.
Parameter values: $\lambda_h + \mu_h = 1$, $\frac{1-p}{p} = \frac{1}{4}$, $\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{1}{2}$

The rate at which soft types are leaked also influences the comparative statics on $\lambda_h + \mu_h$, the effective rate at which hard types screen out. When hard types screen out quickly, soft types learn more quickly and delay less. Nonetheless, this reduced delay turns out to backfire when λ_s is high. Intuitively, this is driven by *rational impatience*: the increase in $\lambda_h + \mu_h$ leads soft types to believe that most hard types have probably already screened out and that it is safe to play hardball against their opponent. While this is a good assumption to make in the limit as $\lambda_h + \mu_h \to \infty$, it does not hold at low levels of $\lambda_h + \mu_h$ where a hard type has not necessarily yet screened out. This effect is compounded when λ_s is high, causing soft types to fear being outed to a soft opponent and being put at a disadvantage. When λ_s is low, the probability of miscoordination decreases in $\lambda_h + \mu_h$. As λ_s increases, nonmonotonicity becomes increasingly pronounced (Figure 5).

Remark 2 ($\lambda_s = 0$ welfare benchmark)**.** *When $\lambda_s = 0$, welfare is increasing in $\lambda_h + \mu_h$, and is decreasing in μ_s.*

These observations highlight how information leaks mediate the effects of commitment opportunities by forcing players to hedge against the probability that they will be caught in the $\overline{T}^*$ case. While higher λ_s is, all else equal, good for welfare, it also means that players are more likely to miscoordinate at intermediate values of μ_s and low-to-intermediate values of $\lambda_h + \mu_h$. By contrast, when λ_s is low, players are most likely to miscoordinate when μ_s is high and $\lambda_h + \mu_h$ is very low – in other words, when soft types can employ hardball tactics frequently and hard types screen out very slowly.

6 Asymmetric setting

In this section, I generalize the analysis to accommodate player- and type- specific rate parameters, beliefs, and utilities. This serves two functions. First, it generalizes the equilibrium statement and provides a clear geometric argument for the uniqueness of the equilibrium in threshold strategies. Second, it provides theoretical depth for the comparative statics described in the symmetric case. Comparative statics previously applied the change of one parameter equally across (the same type of) both players. By disambiguating this effect, I can more clearly separate the effect of one player- and type-specific change on players' best responses. This provides a clearer explanation for the apparently straightforward and ambiguous comparative statics on information leaks and commitment opportunities, respectively. While the purpose of this section is primarily theoretical, it equips the model to answer realistic questions. For instance, a media or legal investigation that aims to leak one political party usually does not threaten the other. Similarly, groups may have different policy priorities, resulting in different relative desirability parameters. The equilibrium consequences of such cases become clear in the asymmetric setting.

6.1 Equilibrium and uniqueness

I retain the setup from Section 3, while relaxing the symmetry assumptions. I allow $u_a^\theta(X) \neq u_b^\theta(X)$ for $\theta \in \{s, h\}$ and generic policy X, and allow $p_a \neq p_b$. Rate parameters are now specific to types *and* players: $\lambda_a^s \neq \lambda_b^s, \lambda_a^h \neq \lambda_b^h$, and likewise for μ parameters. I explicitly notate a soft type of player i's relative desirability as follows:

$$RD_i := \frac{u_i(X_i) - u_i(X_j)}{u_i(X_j) - u_i(SQ)}$$

where X_i denotes i's most-preferred good and X_j denotes j's most-preferred good.[4] The only assumption on parameters that I retain is $\lambda_s^i + \mu_s^i \leq \lambda_h^i + \mu_h^i$. Gameplay proceeds identically to before. In this setup, there exists a unique threshold equilibrium:

Proposition 4 (Uniqueness of the best response correspondence)**.** *Assume that a soft type of player $i = a, b$ is playing a threshold strategy, that is, conditional upon being at a history in which no commitment has been made and no groups type has been leaked, there exists T_i^* such that i will commit to its preferred alternative if and only if $t > T_i^*$. If j uses the threshold time T_j^*, then $T_i^* : \mathbb{R}^+ \to \mathbb{R}^+$ is given by*

$$T_i^*(T_j^*) = \begin{cases} \max\left\{ \frac{1}{\lambda_s^j - (\lambda_h^j + \mu_h^j) - \mu_s^i} \left[\ln\left(RD_i \frac{\mu_s^j}{\mu_s^i + \mu_s^j} \frac{1-p_i}{p_i} \frac{(\lambda_h^j + \mu_h^j) + \mu_s^i}{(\lambda_h^j + \mu_h^j)} \right) - \mu_s^i T_j^* \right], 0 \right\} & \textit{if } T_i < T_j^* \\ \max\left\{ \frac{1}{\lambda_s^j + \mu_s^j - (\lambda_s^j + \mu_s^j)} \left[\ln\left(RD_i \frac{\mu_s^j}{\mu_s^i + \mu_s^j} \frac{1-p_i}{p_i} \frac{(\lambda_h^j + \mu_h^j) + \mu_s^i}{(\lambda_h^j + \mu_h^j)} \right) + \mu_s^j T_j^* \right], 0 \right\} & \textit{if } T_i > T_j^* \end{cases} \tag{6}$$

Furthermore, $T_i^*(T_j^*)$ and $T_j^*(T_i^*)$ are continuous and have a unique point of intersection, which determines optimal delay for each player.

The V-shape of the best response function, shown in Figure 6, suggests why players collectively suffer most in the symmetric case. In the symmetric case, players delayed an equal amount of time, causing their best responses to intersect on their (identical) kink: T^*. In the asymmetric case, best response functions are still piecewise and kink on the 45-degree line, but may not kink and intersect at the same point. Figure 6 depicts best responses intersecting where $T_a^* < T_b^*$. This means that b cannot learn from commitment behavior after T_a^*. Analogously to the argument for $\overline{T}^* > T^*$ in the symmetric case, b's slower learning translates into longer delay.

[4] Note that varying relative desirability is equivalent to varying the payoff that players receive in cases where negotiation fails, since what matters in players' strategic calculus is the value of the "status quo" that can be received in the future, not the value of the one received in the past.

Whenever a soft type of the opponent becomes willing to commit earlier, the decisionmaker delays longer. The opponent with the earlier threshold, however, does not delay as long as in symmetric case. They delay *more*. They have the flexibility to extend delay slightly, reducing their chances of preempting a hard type mistakenly without overly increasing the chance their lagging opponent will preempt them. Neither player delays less than in the symmetric case. The result is particularly stark in cases like this one where only one player's best response responds to a change to a parameter. From Proposition 4, we know that p_b, b's prior about a, only appears in b's best response. Hence, b's best response function (in blue) is shifted while a's (in red) remains unchanged from the symmetric case depicted in Figure 2.

The asymmetric setting also offers insights to why comparative statics on λ_s tended to be simpler than on other rate parameters: Increasing λ_s^i, the rate at which player i is leaked, shifts j's best response without shifting i's (see Proposition 4). As a result, increasing λ_s^i must increase delay for both groups because, like changing one group's prior, it holds one group's best response constant. Returning to the symmetric case, we can see that this effect is only compounded by increasing both λ_s^a and λ_s^b. This shifts *both* best responses outwards, compounding the delaying effect on both players.

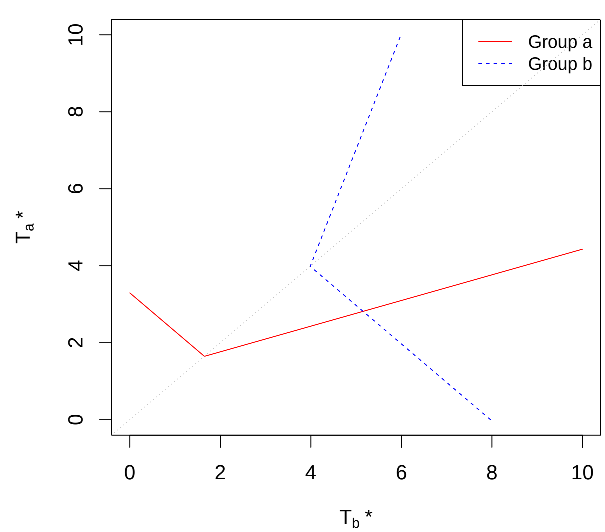

(a) Best response correspondences when a is more likely than b to be a hard type ($p_b = 0.95 > p_a = 0.8$)

Figure 6. Best responses in an asymmetric setting. The 45-degree line indicating when $T_a^* = T_b^*$ is marked in light grey. Other parameter values are preserved from Figure 2.

Unlike λ_s^i, changing μ_s^i or $\lambda_h^i + \mu_h^i$ shifts best response functions for *both* players. While analytic characterizations are more difficult, comparative statics done in simulation confirm basic intuitions (see Figure A3). Increasing μ_s^i, the average rate at which a soft type of player i receives commitment opportunities, leads both players to delay less, with j delaying the least. Increasing $\lambda_h^i + \mu_h^i$ also leads to less delay, especially by player j, who can make faster inferences about their opponent's type. These are consistent with the symmetric case, where T^* was decreasing in μ_s as well as in $\lambda_h + \mu_h$.

7 Discussion

This paper develops a model of how groups with competing interests choose to employ hardball tactics against each other. I show that groups who are unwilling to compromise on their interests play hardball quickly, while groups who are willing to compromise delay usage of these tactics. Because groups commit to rather than concede from their demands, the dynamic learning process is reversed from reputational bargaining: As time passes, a group who has yet to act is more likely in its opponent's eyes to be a compromising "soft" type than an obstinate "hard" type. While uncompromising hard types behave mechanically, soft types' strategic delay is responsive to the underlying parameters of the game, including the frequency with which groups can act, prior beliefs about types, and the intensity of preferences. The main source of inefficiency in the model is the probability of avoidable miscoordination, which arises due to the stochastic and unobserved arrival of opportunities to play hardball in continuous time. The possibility of avoidable miscoordination highlights the importance of credible actions that put an end to potential compromises, rather than "cheap talk" demands that allow players to take aggressive-seeming negotiation positions without costly skin in the game. Even if such communication were to be incorporated into the model, players would remain reluctant to play hardball unless sufficiently confident in their opponent's willingness to compromise.

Analysis of uninformed soft types' decisions coalesce into two major themes. First, any form of asymmetry, either induced endoegnously through leaks or exogenously through asymmetric parameter values, impedes one player's learning, causing them to delay playing hardball for longer. Second, equilibrium behavior connects to welfare through the risk of avoidable miscoordination, which is mollified by factors that prolong delay. Relatedly, while increasing the likelihood of leaks had analytically simple effects on equilibrium behavior and welfare, the likelihood of leaks mediated the effects of faster or slower commitment opportunities. The root cause of this is players' *ex ante* uncertainty about whether they will find themselves at an informational disadvantage mid-game. These comparative statics on information leaks and commitment opportunities are absent from most bargaining games and games of preemption, and address real-life questions. For instance, a group is revealed to be willing to compromise, we should expect that group to both abstain longer from hardball tactics, but also that an impasse is less likely. If factions are able to have more frequent meetings which translate into commitment opportunities, negotiations may be shorter, but are also more likely to stalemate.

I now briefly revisit a key assumption, namely that $\lambda_s + \mu_s \leq \lambda_h + \mu_h$. This assumption is necessary to obtain a threshold equilibrium wherein soft types never commit before the threshold and always commit after the threshold. If the assumption is reversed, the logic of soft types' decision prior to the threshold is unchanged (since hard types still want to commit quickly), but after the threshold, *soft types will screen out faster than hard types.* Therefore, if no commitments are observed after the threshold, players think that *an opponent who has not committed is more likely to be a hard type.* In short, undoing this assumption means that learning reverses after the threshold. In order to equalize opponents' beliefs at the threshold without the assumption, soft types play a mixed strategy between committing and not committing. Since this does not present a substantial qualitative difference from the results obtained under the original assumption, I exclude it for analytical clarity.

The model is flexible to a number of extensions, for instance, imposing negotiation deadlines. Suppose, for example, that there is a known time T_D past which, if either group has failed to make a commitment, the status quo is automatically instated. The effect of this is to add an additional cost to uninformed soft types' cost of waiting, corresponding to the probability that a commitment opportunity will not arrive prior to the deadline $(1 - e^{-\mu_s(T_D - t)})$. If the deadline is sufficiently early, soft types would be strongly incentivized towards preemption, making learning difficult (since players of both types act quickly). If it is relatively late, the threshold equilibrium would be largely unchanged, although it should shift slightly earlier due to the extra weight on the costs of delay.

Related questions remain open, for instance, that of endogenous information acquisition. If a group can pay a cost to do "opposition research" by increasing λ_s or λ_h of its opponent, when would it pay this cost? How would this affect each group's expected payoff? I leave this to future work expanding on the model. The basic framework advanced in this model is not limited to the application of factional competition in political institutions. For instance, it could also be used to model the logic of sales or employment negotiations. In practice, both seller and buyer must agree to the same contract for negotiations to go forward, with both parties having enough leverage to set a price above which they will not buy, or below which they will not sell.[5] Alternatively, in hiring decisions that are made by committee, stakeholders with different priorities must frequently agree on a single candidate. If one group of stakeholders strongly prefers candidate a and another strongly prefers candidate b, each faction must decide when to "pull the trigger" based on their perception of the other side's willingness to compromise on their choice.

[5]This is reminiscent of a passage in Schelling (1956), which employs the example of a house seller and buyer: "Suppose the buyer could make an irrevocable and enforceable bet with some third party, duly recorded and certified, according to which he would pay for the house no more than $16,000 (...) **The seller can take it or leave it. This example demonstrates that if the buyer can accept an irrevocable commitment, in a way that is unambiguously visible to the seller, he can squeeze the range of indeterminacy down to the point most favorable to him.**" (Emphasis added.)

References

Abreu, Dilip and Faruk Gul. 2000. "Bargaining and Reputation." *Econometrica* 68(1):85–117.

Ambrus, Attila and Shih En Lu. 2015. "A continuous-time model of multilateral bargaining." *American Economic Journal: Microeconomics* 7(1):208–249.

Bade, Rachel, Josh Dawsey and Josh Haberkorn. 2017. "How a secret Freedom Caucus pact brought down Obamacare repeal." *Politico* . Available at `https://www.politico.com/story/2017/03/trump-freedom-caucus-obamacare-repeal-replace-secret-pact-236507`.

Calcagno, Riccardo, Yuichiro Kamada, Stefano Lovo and Takuo Sugaya. 2014. "Asynchronicity and coordination in common and opposing interest games." *Theoretical Economics* 9(2):409–434.

Calvo, Guillermo A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics* 12(3):383–398.

Cameron, Charles and Nolan McCarty. 2004. "Models of vetoes and veto bargaining." *Annu. Rev. Polit. Sci.* 7:409–435.

Cox, Gary W and Mathew D McCubbins. 2005. *Setting the agenda: Responsible party government in the US House of Representatives.* Cambridge University Press.

De Mesquita, Bruce Bueno, Alastair Smith, Randolph M Siverson and James D Morrow. 2005. *The logic of political survival.* MIT press.

Dendrinou, Viktoria and Gabriele Steinhauser. 2016. "Eurozone and IMF Strike Deal on Greek Debt." *The Wall Street Journal* . Available at `https://www.wsj.com/articles/eurozone-and-imf-strike-deal-on-greek-debt-1464137992`.

Fearon, James D. 1994. "Domestic political audiences and the escalation of international disputes." *American Political Science Review* 88(3):577–592.

Fey, Mark and Brenton Kenkel. 2021. "Is an ultimatum the last word on crisis bargaining?" *The Journal of Politics* 83(1):87–102.

Fudenberg, Drew and Jean Tirole. 1986. "A Theory of Exit in Duopoly." *Econometrica* 54(4):943–960.

Green, Matthew. 2019. *Legislative Hardball: The House Freedom Caucus and the Power of Threat-Making in Congress.* Cambridge University Press.

Horowicz, Alejandro. 2009. "Ezeiza Protest and Massacre, 1973." *The International Encyclopedia of Revolution and Protest* pp. 1–3.

Invernizzi, Giovanna M. 2023. "Antagonistic cooperation: Factional competition in the shadow of elections." *American Journal of Political Science* 67(2):426–439.

Izzo, Federica. 2020. "With friends like these, who needs enemies?" *Working paper* . Available at `https://www.federicaizzo.com/pdf/WFLT10K.pdf`.

James, Harold. 2024. The IMF and the European Debt Crisis. In *The IMF and the European Debt Crisis*. International Monetary Fund.

Kamada, Yuichiro and Michihiro Kandori. 2020. "Revision games." *Econometrica* 88(4):1599–1630.

Kamada, Yuichiro and Takuo Sugaya. 2020. "Optimal timing of policy announcements in dynamic election campaigns." *The Quarterly Journal of Economics* 135(3):1725–1797.

Kiewiet, D Roderick and Mathew D McCubbins. 1991. *The logic of delegation*. University of Chicago Press.

Krehbiel, Keith. 2010. *Pivotal politics: A theory of US lawmaking*. University of Chicago Press.

Meng, Anne, Jack Paine and Robert Powell. 2023. "Authoritarian Power Sharing: Concepts, Mechanisms, and Strategies." *Annual Review of Political Science* 26:153–173.

O'Donnell, Guillermo and Philippe C Schmitter. 1986. *Transitions from Authoritarian Rule: Tentative Conclusions about Uncertain Democracies*. The Johns Hopkins Universit Press.

Rubin, Ruth Bloch. 2017. *Building the bloc: Intraparty organization in the US Congress*. Cambridge University Press.

Schelling, Thomas C. 1956. "An essay on bargaining." *The American Economic Review* 46(3):281–306.

Svolik, Milan W. 2009. "Power sharing and leadership dynamics in authoritarian regimes." *American Journal of Political Science* 53(2):477–494.

Walker, Marcus and Nektaria Stamouli. 2016. "Leaked Call Shows Continued Strains Between Greece and Its Creditors." *The Wall Street Journal* . Available at `https://www.wsj.com/articles/leaked-call-shows-continued-strains-between-greece-and-its-creditors-1459617458`.

Weeds, Helen. 2002. "Strategic delay in a real options model of R&D competition." *The Review of Economic Studies* 69(3):729–747.

Online Appendix

A Proofs for equilibrium in the symmetric setting

A.1 Derivation of $T^*, \overline{T}^*$

Derivation of T^*. At T^*, a soft type of group a must be indifferent between committing to its preferred policy and waiting. (The derivation is analogous for a soft type or group b). Then, a's expected utility of committing to A at T^*, conditional upon neither type being revealed and no commitments occurring by T^*, is

$$\Big[pe^{(-\mu_h-\lambda_h)T^*}\Big]u_a(SQ) + \Big[(1-p)e^{-\lambda_s T^*}\Big]u_a(A) \tag{A1}$$

To derive the continuation value of waiting at T^*, consider the possible subsequent events after a waits. Suppose first that b is hard. Then, two cases are possible: (1) μ_h arrives, in which case b commits to B, or λ_h arrives, in which case a knows that implementing A is not possible. Either way, a soft type of a prefers B to SQ, so a commits to B as soon as it receives a commitment opportunity. (2) μ_s arrives. Assuming that a is playing the proposed equilibrium strategy, a commits to policy A. However, since b prefers SQ to A, the outcome is SQ. Suppose now that b is soft. Then, whichever group receives the first commitment opportunity after T^* commits to their preferred policy. Since μ_s is the same for both players, these occur with equal probability. Whichever group receives the first commitment opportunity can implement its preferred policy. Thus, the continuation value of waiting is

$$\Big[pe^{(-\mu_h-\lambda_h)T^*}\Big(\frac{1}{\mu_s+\mu_h+\lambda_h}\Big)\Big]\Big(\mu_s u_a(SQ) + (\mu_h+\lambda_h)u_a(B)\Big) + \Big[\frac{1-p}{2}e^{-\lambda_s T^*}\Big]\Big(u_a(A)+u_a(B)\Big) \tag{A2}$$

Setting (A1) equal to (A2) and rearranging terms, I obtain a's indifference condition at T^*:

$$\Big[pe^{(-\mu_h-\lambda_h)T^*}\Big(\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\Big)\Big]\Big(u_a(B)-u_a(SQ)\Big) = \Big[\frac{1-p}{2}e^{-\lambda_s T^*}\Big]\Big(u_a(A)-u_a(B)\Big) \tag{A3}$$

As the probabilities used in the calculation of both these expressions are technically conditional probabilities (conditioned upon neither player's type being revealed before the current time), we are obliged to divide each probability by the sum total of all the probabilities used in the calculation of expected utility in order for probabilities to sum to 1. The sum of the associated probabilities of committing to A is, trivially

$$pe^{(-\mu_h-\lambda_h)T^*} + (1-p)e^{-\lambda_s T^*}$$

The sum of the associated probabilities of waiting should be equal this exactly, as they should theoretically both yield the total probability of no commitments and no preference revelations before T^*. Indeed, the sum of all probabilities associated with waiting is

$$
\begin{aligned}
&pe^{(-\mu_h-\lambda_h)T^*}\frac{\mu_s+\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}+(1-p)e^{-\lambda_s T^*}\frac{\mu_s+\mu_s}{\mu_s+\mu_s}\\
=&pe^{(-\mu_h-\lambda_h)T^*}+(1-p)e^{-\lambda_s T^*}
\end{aligned}
$$

which confirms that the normalization factors are then equal. We then note that once we set the expected values equal, each term will be divided by the same normalization factor, so they will cancel.

Rearrange terms to isolate T^*:

$$T^* = \frac{1}{\lambda_s-(\lambda_h+\mu_h)}\left[\ln\left(\frac{1-p}{2p}\,\frac{(\lambda_h+\mu_h)+\mu_s}{(\lambda_h+\mu_h)}\,\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)}\right)\right]$$

Derivation of $\overline{T}^*$. Following similar logic as before, a's expected utility of committing to A at time $\overline{T}^*$ is

$$\left(pe^{(-\mu_h-\lambda_h)\overline{T}^*}\right)u_a(SQ)+\left((1-p)e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\right)u_a(A) \tag{A4}$$

The main difference with the previous derivation is that the opponent is now attempting to commit during the interval $\overline{T}^*-\bar{t}$. a's continuation value of waiting at $\overline{T}^*$ is

$$\left[\frac{1-p}{2}e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\right]\Big(u_a(A)-u_a(B)\Big)=\left[pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h+\mu_s}\right]\Big(u_a(B)-u_a(SQ)\Big) \tag{A5}$$

Setting (A4) equal to (A5) and rearranging terms, I obtain a's indifference condition at $\overline{T}^*$:

$$\left[\frac{1-p}{2}e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\right]\Big(u_a(A)-u_a(B)\Big)=\left[pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h+\mu_s}\right]\Big(u_a(B)-u_a(SQ)\Big) \tag{A6}$$

Similarly to before, normalization terms cancel out. Rearrange terms to isolate $\overline{T}^*$, which is a linear function of $\bar{t}$:

$$\overline{T}^* = \frac{1}{\lambda_s+\mu_s-(\lambda_h+\mu_h)}\left[\ln\left(\frac{1-p}{2p}\,\frac{(\lambda_h+\mu_h)+\mu_s}{(\lambda_h+\mu_h)}\,\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)}\right)+\mu_s\bar{t}\,\right] \tag{A7}$$

For completeness, note that while these expression can become negative, players cannot wait a negative period of time. The player's threshold time for action is therefore $\max\{0,T^*\}$. The same applies to $\overline{T}^*$.

A.2 Proof of Propositions 1 and 2

I will show that conditional on no commitments and no leaks, a soft type prefers to wait before T^* and prefers to commit to its preferred policy after T^*, and that conditional on its type being revealed at $\bar{t}$, a soft type prefers to wait before $\overline{T}^*$ and prefers to commit to its preferred policy after $\overline{T}^*$.

A.2.1 Proposition 1 (T^* case)

Claim: **When $t < T^*$, a strictly prefers to wait.**

For all $t < T^*$, a soft type of player a's expected utility of committing to A is:

$$pe^{(-\mu_h-\lambda_h)t}u_a(SQ) + (1-p)e^{-\lambda_s t}u_a(A) \tag{A8}$$

I now derive a soft type of group a's continuation value of waiting at t. I proceed in cases. The possible cases when b is a hard type are:

1. μ_h or λ_h arrives before T^*. Policy B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)T^*}$
2. Either μ_h or λ_h arrive before μ_s but after T^*. Policy B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)T^*}\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h+\mu_s}$
3. Neither μ_h nor λ_h arrive before T^*, and μ_s arrives before μ_h and λ_h after T^*. Policy SQ is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)T^*}\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}$.

The possible cases when b is a soft type are:

4. λ_s arrives for either group at some $\bar{t}$ before T^*, and the other group receives a commitment opportunity before $\overline{T}^*(\bar{t})$, and commits to their preferred policy, which is implemented. This occurs with probability

   $$(1-p)e^{-\lambda_s t}\int_{\tilde{t}=t}^{T^*} e^{-2\lambda_s(\tilde{t}-t)}(\lambda_s)e^{-\mu_s[\overline{T}^*(\bar{t})-\tilde{t}]}d\tilde{t}$$

   Note that $e^{-2\lambda_s(\tilde{t}-t)}(\lambda_s)$ is the *instantaneous* probability of an arrival of λ_s at any instant $\tilde{t}$. I then multiply this by the probability that, conditional upon this arrival happening at some $\tilde{t}$, μ_s arrives between $\tilde{t}$ and $\overline{T}^*(\tilde{t})$. Evaluating this expression yields

$$\left(\frac{\lambda_s \frac{(1-p)^2}{p} \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} \frac{\lambda_h+\mu_h+\mu_s}{\mu_h+\lambda_h}}{-2\lambda_s + \mu_s - \frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h}}\right) e^{\lambda_s t - \frac{\mu_s}{\lambda_s+\mu_s-\lambda_h-\mu_h}} \left[e^{(-2\lambda_s-\mu_s-\frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h})T^*} - e^{(-2\lambda_s-\mu_s-\frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h})t}\right]$$

$$\equiv P_4$$

5. λ_s arrives for either group at some $\bar{t}$ before T^*, and the other group does not get a commitment opportunity before $\overline{T}^*(\bar{t})$. After $\overline{T}^*$, whichever group gets the first commitment opportunity is able to implement their preferred policy. This occurs with probability

$$(1-p)e^{-\lambda_s t} \int_{\tilde{t}=t}^{T^*} e^{-2\lambda_s(\tilde{t}-t)}(\lambda_s)\left[1 - e^{-\mu_s[\overline{T}^*(\bar{t})-\tilde{t}]}\right] d\tilde{t}$$

Evaluating this expression yields

$$\frac{-(1-p)}{2}\left[e^{\lambda_s(t-2T^*)} - e^{-\lambda_s t}\right] - \left(\frac{\lambda_s \frac{(1-p)^2}{p} \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} \frac{\lambda_h+\mu_h+\mu_s}{\mu_h+\lambda_h}}{-2\lambda_s + \mu_s - \frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h}}\right) e^{\lambda_s t - \frac{\mu_s}{\lambda_s+\mu_s-\lambda_h-\mu_h}} \left[e^{(-2\lambda_s-\mu_s-\frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h})T^*} - e^{(-2\lambda_s-\mu_s-\frac{\mu_s^2}{\lambda_s+\mu_s-\lambda_h-\mu_h})t}\right]$$

$$= \frac{-(1-p)}{2}\left[e^{\lambda_s(t-2T^*)} - e^{-\lambda_s t}\right] - P_4$$

6. There are no arrivals of λ_s before T^* Whichever group receives the first commitment opportunity after T^* commits to their preferred policy, which is implemented. This occurs with probability

$$\frac{(1-p)}{2} e^{\lambda_s t - 2\lambda_s T^*}$$

Therefore, the continuation value of waiting at t is

$$\left(pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)T^*} \frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\right) u_a(B) + \left(pe^{(-\mu_h-\lambda_h)T^*} \frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\right) u_a(SQ) + \frac{(1-p)}{2} e^{-\lambda_s t}\left(\frac{u_a(A)+u_a(B)}{2}\right) \tag{A9}$$

In order for waiting to be optimal, we must have (A8) < (A9), which simplifies to

$$\left(pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)T^*} \frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\right)\Big(u_a(B) - u_a(SQ)\Big) > \frac{1-p}{2} e^{-\lambda_s t}\Big(u_a(A) - u_a(B)\Big) \tag{A10}$$

Because $t < T^*$, it holds that

$$pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)T^*}\frac{\mu_s}{\lambda_h+\mu_h+\mu_s} > pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)t}\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}$$
$$= pe^{(-\mu_h-\lambda_h)t}\frac{\mu_h+\lambda_h}{\lambda_h+\mu_h+\mu_s}$$

It must also be that

$$\left(pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)T^*}\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\right)\Big(u_a(B)-u_a(SQ)\Big)$$
$$>\left(pe^{(-\mu_h-\lambda_h)t}\frac{\mu_h+\lambda_h}{\lambda_h+\mu_h+\mu_s}\right)\Big(u_a(B)-u_a(SQ)\Big)$$

Therefore, it suffices to prove that

$$\left(pe^{(-\mu_h-\lambda_h)t}\frac{\mu_h+\lambda_h}{\lambda_h+\mu_h+\mu_s}\right)\Big(u_a(B)-u_a(SQ)\Big) > \left(\frac{1-p}{2}e^{-\lambda_s t}\right)\Big(u_a(A)-u_a(B)\Big) \quad \text{(A11)}$$

Recall the indifference condition for T^* derived earlier:

$$\left(pe^{(-\mu_h-\lambda_h)T^*}\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\right)\Big(u_a(B)-u_a(SQ)\Big) = \left(\frac{1-p}{2}e^{-\lambda_s T^*}\right)\Big(u_a(A)-u_a(B)\Big) \quad \text{(A3)}$$

Notice that

$$\frac{\left[pe^{(-\mu_h-\lambda_h)t}\frac{\mu_h+\lambda_h}{\lambda_h+\mu_h+\mu_s}\right]\Big(u_a(B)-u_a(SQ)\Big)}{\left[\frac{1-p}{2}e^{-\lambda_s t}\right]\Big(u_a(A)-u_a(B)\Big)}$$
$$=\frac{\left[pe^{(-\mu_h-\lambda_h)T^*}\left(\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\right)\right]\Big(u_a(B)-u_a(SQ)\Big)}{\left[\frac{1-p}{2}e^{-\lambda_s T^*}\right]\Big(u_a(A)-u_a(B)\Big)}\cdot e^{(-\mu_h-\lambda_h+\lambda_s)(t-T^*)}$$

Consider the factor at the end of the expression, $e^{(-\mu_h-\lambda_h+\lambda_s)(t-T^*)}$. Since $(-\mu_h - \lambda_h + \lambda_s) < 0$ and $(t - T^*) < 0$, we have $(-\mu_h - \lambda_h + \lambda_s)(t - T^*) > 0$ and therefore $e^{(-\mu_h-\lambda_h+\lambda_s)(t-T^*)} > 1$. Therefore the left-hand side of (A11) is greater than the right-hand side and the inequality is true. This concludes the proof that when $t < T^*$, a strictly prefers to wait.

Claim: **When $t > T^*$, a strictly prefers to commit to A.**

Let $\epsilon > 0$. At time $T^* + \epsilon$, a soft type of player a's expected utility of committing to A is

$$\left[pe^{(T^*+\epsilon)(-\mu_h-\lambda_h)}\right]u_a(SQ) + \left[(1-p)e^{-\mu_s\epsilon-\lambda_s(T^*+\epsilon)}\right]u_a(A) \quad \text{(A12)}$$

I now derive a soft type of group a's continuation value of waiting at $T^* + \epsilon$. I proceed in cases:

1. b is a hard type. μ_h or λ_h arrives first. B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*+\epsilon)}\left(\frac{\lambda_h+\mu_h}{\mu_h+\lambda_h+\mu_s}\right)$
2. b is a hard type. μ_s arrives first. SQ remains in place. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*+\epsilon)}\left(\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right)$
3. b is a soft type. In this case, the first player to get a commitment opportunity after $T^*+\epsilon$ is able to implement their preferred policy. Either A or B is implemented, each occurs with probability $\frac{(1-p)}{2}e^{(-\lambda_s-\mu_s)(\overline{T}^*+\epsilon)-\mu_s\bar{t}}$

Committing to A is preferable to waiting if

$$\left[pe^{(-\mu_h-\lambda_h)(T^*+\epsilon)}\Big(\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\Big)\right]\Big(u_a(B)-u_a(SQ)\Big) < \left[\frac{1-p}{2}e^{-\mu_s\epsilon-\lambda_s(T^*+\epsilon)}\right]\Big(u_a(A)-u_a(B)\Big) \tag{A13}$$

Recall equation A3 that describes indifference at T^*:

$$\left[pe^{(-\mu_h-\lambda_h)T^*}\Big(\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\Big)\right]\Big(u_a(B)-u_a(SQ)\Big) = \left[\frac{1-p}{2}e^{-\lambda_s T^*}\right]\Big(u_a(A)-u_a(B)\Big)$$

Note that

$$\frac{pe^{(-\mu_h-\lambda_h)(T^*+\epsilon)}\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\Big(u_a(B)-u_a(SQ)\Big)}{\frac{1-p}{2}e^{-\mu_s\epsilon-\lambda_s(T^*+\epsilon)\Big(u_a(A)-u_a(B)\Big)}\Big(u_a(A)-u_a(B)\Big)}$$
$$=\frac{pe^{(-\mu_h-\lambda_h)T^*}\left(\frac{\mu_h+\lambda_h}{\mu_s+\mu_h+\lambda_h}\right)\Big(u_a(B)-u_a(SQ)\Big)}{\frac{1-p}{2}e^{-\lambda_s T^*}\Big(u_a(A)-u_a(B)\Big)}\cdot e^{(\mu_s+\lambda_s-(\lambda_h+\mu_h))\epsilon}$$

By assumption that $\lambda_s+\mu_s<\lambda_h+\mu_h$, we have $e^{(\mu_s+\lambda_s-(\lambda_h+\mu_h))\epsilon}<1$. Therefore inequality (A13) holds. This concludes the proof that when $t>T^*$, a strictly prefers to commit to A.

A.2.2 Proposition 2 ($\overline{T}^*$ case)

Claim: **Suppose a's type was revealed at some time $\bar{t}$. Then a strictly prefers to wait at any $t<\overline{T}^*$.**

Let $t>\bar{t}$. a's expected utility of committing to A at time t is

$$(1-p)e^{-\mu_s(t-\bar{t})-\lambda_s t}u_a(A)+pe^{(-\mu_h-\lambda_h)t}u_a(SQ)$$

I now derive a soft type of group a's continuation value of waiting at t.

If b is a hard type,

1. μ_h or λ_h arrives before $\overline{T}^*(\bar{t})$. B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)t} - pe^{(-\mu_h-\lambda_h)\overline{T}^*(\bar{t})}$

2. μ_h or λ_h do not arrive between t and $\overline{T}^*(\bar{t})$, but arrive before μ_s after $\overline{T}^*(\bar{t})$. B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*(\bar{t})}\frac{\mu_h+\lambda_h}{\mu_h+\lambda_h+\mu_s}$

3. μ_h or λ_h do not arrive between t and $\overline{T}^*(\bar{t})$, and μ_s arrives first after $\overline{T}^*(\bar{t})$. SQ remains in place. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*(\bar{t}))}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}$

If b is a soft type,

1. μ_s arrives for b between $\bar{t}$ and $\overline{T}^*(\bar{t})$. B is implemented. This occurs with probability

$$\begin{aligned}&(1-p)e^{-\mu_s(t-\bar{t})-\lambda_s t}\int_{\tilde{t}=t}^{\overline{T}^*} e^{-\lambda_s(\tilde{t}-t)-\mu_s(\tilde{t}-t)}(\mu_s)d\tilde{t}\\ =&(1-p)e^{\mu_s\bar{t}}\frac{\mu_s}{-\mu_s-\lambda_s}\left[e^{(-\lambda_s-\mu_s)\overline{T}^*}-e^{(-\lambda_s-\mu_s)t}\right]\end{aligned}$$

2. λ_s arrives for b before $\overline{T}^*$. Both groups are fully informed, and the first that receives a commitment opportunity implements their preferred policy. Either A or B is implemented. Each sub-case occurs with probability

$$\begin{aligned}&(1-p)e^{-\mu_s(t-\bar{t})-\lambda_s t}\int_{\tilde{t}=t}^{\overline{T}^*} e^{-\lambda_s(\tilde{t}-t)-\mu_s(\tilde{t}-t)}(\lambda_s)d\tilde{t}\\ =&(1-p)e^{\mu_s\bar{t}}\frac{\lambda_s}{-\mu_s-\lambda_s}\left[e^{(-\lambda_s-\mu_s)\overline{T}^*}-e^{(-\lambda_s-\mu_s)t}\right]\end{aligned}$$

3. Neither μ_s nor λ_s arrive for b before $\overline{T}^*$. The first group that receives a commitment opportunity implements their preferred policy. Either A or B is implemented. Each sub-case occurs with probability

$$\begin{aligned}&\frac{1-p}{2}e^{-\mu_s(t-\bar{t})-\lambda_s t}e^{-\lambda_s(\overline{T}^*-t)-\mu_s(T-1-t)}\\ &=\frac{1-p}{2}e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\end{aligned}$$

Waiting is preferable to committing to A if

$$\begin{aligned}&\left[pe^{(-\mu_h-\lambda_h)t}-pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right](u_a(B)-u_a(SQ))\\ &+(1-p)e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\left(\frac{u_a(A)+u_a(B)}{2}+\frac{\lambda_s(u_a(A)+u_a(B))+\mu_s u_a(B)}{-\lambda_s-\mu_s}\right)\\ >&(1-p)e^{-\lambda_s t-\mu_s(t-\bar{t})}\left[u_a(A)+\frac{\lambda_s(u_a(A)+u_a(B))+\mu_s u_a(B)}{-\lambda_s-\mu_s}\right]\end{aligned} \tag{A14}$$

Which holds if the following holds:

$$\begin{aligned}&\left[pe^{(-\mu_h-\lambda_h)t}-pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right](u_a(B)-u_a(SQ))\\&+(1-p)e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\left(\frac{u_a(A)+u_a(B)}{2}+\frac{\lambda_s(u_a(A)+u_a(B))+\mu_s u_a(B)}{-\lambda_s-\mu_s}\right)\\>&(1-p)e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\left[u_a(A)+\frac{\lambda_s\big(u_a(A)+u_a(B)\big)+\mu_s u_a(B)}{-\lambda_s-\mu_s}\right]\end{aligned} \tag{A15}$$

Which simplifies to

$$\left[\frac{1-p}{2}e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\right](u_a(A)-u_a(B))<\left[pe^{(-\mu_h-\lambda_h)t}-pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right](u_a(B)-u_a(SQ)) \tag{A16}$$

Note that the right-hand side is greater than

$$\begin{aligned}&\left[pe^{(-\mu_h-\lambda_h)\overline{T}^*}-pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right](u_a(B)-u_a(SQ))\\=&\left[pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h+\mu_s}\right]\Big(u_a(B)-u_a(SQ)\Big)\end{aligned}$$

which is the right-hand side of the following equation, which was the condition for a to be indifferent at $\overline{T}^*$:

$$\left[\frac{1-p}{2}e^{-\lambda_s\overline{T}^*-\mu_s(\overline{T}^*-\bar{t})}\right]\Big(u_a(A)-u_a(B)\Big)=\left[pe^{(-\mu_h-\lambda_h)\overline{T}^*}\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h+\mu_s}\right]\Big(u_a(B)-u_a(SQ)\Big) \tag{A6}$$

Since the left-hand side of (A16) is identical to that of equation A6, we must have that the inequality must be true. Thus, a strictly prefers to wait at any $t<\overline{T}^*$.

Claim: **Suppose a's type was revealed at some time $\bar{t}$. Then a strictly prefers to commit to A at any $t>\overline{T}^*$.**

Let $\epsilon>0$. a's expected utility of committing to A is

$$pe^{(-\mu_h-\lambda_h)(\overline{T}^*+\epsilon)}u_a(SQ)+(1-p)e^{(-\lambda_s-\mu_s)(\overline{T}^*+\epsilon)-\mu_s\bar{t}}u_a(A)$$

I now derive a soft type of group a's continuation value of waiting at $\overline{T}^*+\epsilon$. I proceed in cases:

1. b is a hard type. μ_h or λ_h arrives before $\overline{T}^*$. Policy B is implemented. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*+\epsilon)}\left(\frac{\mu_h+\lambda_h}{\mu_h+\lambda_h+\mu_s}\right)$
2. b is a hard type. Neither μ_h nor λ_h arrive before $\overline{T}^*$, and μ_s arrives before μ_h and λ_h after $\overline{T}^*$. SQ remains in place. This occurs with probability $pe^{(-\mu_h-\lambda_h)(\overline{T}^*+\epsilon)}\left(\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\right)$

3. b is a soft type. Then the first player to receive a commitment opportunity after $\overline{T}^* + \epsilon$ is able to implement their preferred alternative. Either A or B is implemented. The probability of each sub-case is $\frac{(1-p)}{2} e^{(-\lambda_s - \mu_s)(\overline{T}^* + \epsilon) - \mu_s \bar{t}}$

a prefers to commit to A at any $\overline{T}^* + \epsilon$ if the following inequality holds:

$$\left[\frac{1-p}{2} e^{(-\lambda_s - \mu_s)(\overline{T}^* + \epsilon) - \mu_s \bar{t}}\right]\Big(u(A) - u(B)\Big) > \left[p e^{(-\mu_h - \lambda_h)(\overline{T}^* + \epsilon)} \frac{\mu_h + \lambda_h}{\mu_h + \lambda_h + \mu_s}\right]\Big(u(B) - u(SQ)\Big)$$

Recall the condition for a to be indifferent at $\overline{T}^*$ (equation A6) was:

$$\left[\frac{1-p}{2} e^{-\lambda_s \overline{T}^* - \mu_s(\overline{T}^* - \bar{t})}\right]\Big(u_a(A) - u_a(B)\Big) = \left[p e^{(-\mu_h - \lambda_h)\overline{T}^*} \frac{\lambda_h + \mu_h}{\lambda_h + \mu_h + \mu_s}\right]\Big(u_a(B) - u_a(SQ)\Big)$$

Note that

$$\frac{\left[p e^{(-\mu_h - \lambda_h)(\overline{T}^* + \epsilon)} \frac{\mu_h + \lambda_h}{\mu_h + \lambda_h + \mu_s}\right]\Big(u(B) - u(SQ)\Big)}{\left[\frac{1-p}{2} e^{(-\lambda_s - \mu_s)(\overline{T}^* + \epsilon) - \mu_s \bar{t}}\right]\Big(u(A) - u(B)\Big)}$$

$$= \frac{\left[p e^{(-\mu_h - \lambda_h)\overline{T}^*} \frac{\lambda_h + \mu_h}{\lambda_h + \mu_h + \mu_s}\right]\Big(u(B) - u(SQ)\Big)}{\left[\frac{1-p}{2} e^{-\lambda_s \overline{T}^* - \mu_s(\overline{T}^* - \bar{t})}\right]\Big(u(A) - u(B)\Big)} \cdot e^{(\lambda_s + \mu_s - (\lambda_h + \mu_h))\epsilon}$$

By assumption that $\lambda_h + \mu_h > \lambda_s + \mu_s$, we have that $e^{(\lambda_s + \mu_s - (\lambda_h + \mu_h))\epsilon}$ and therefore the inequality holds. This concludes the proof that a strictly prefers to commit to A at any $t > \overline{T}^*$.

Proof of Corollary 1 (conditions for no delay). As either $\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \to \infty$ or $\frac{1-p}{p} \to infty$, we must have that $\frac{1-p}{p} P_T \to \infty$, and $\log \infty = \infty$. As this is multiplied by a negative number in both delay terms, and delay must be nonnegative, delay $\to 0$.

A.3 Proofs of comparative statics on rates (Proposition 3)

	λ_s	μ_s	$\mu_h + \lambda_h$
P_T	.	$(+)$	$(-)$
T^*	$(+)$	$(-)$	$(-)$
$\overline{T}^*$	$(+)$	$(+/-)$	$(-)$

Table A1. Summary of comparative statics on λ_s, μ_s, and $\lambda_s + \lambda_h$

A.3.1 P_T

Claim: P_T **is decreasing in** $\lambda_h + \mu_h$**.**

$$\frac{\partial P_T}{\partial \lambda_h + \mu_h} = \frac{1}{2} \frac{u(A) - u(B)}{u(B) - u(SQ)} \frac{-\mu_s}{(\lambda_h + \mu_h)^2}$$

The last term is negative and is multiplied by positive terms, so the derivative is negative.

Claim: P_T **is increasing in** μ_s**.**

$$\frac{\partial P_T}{\partial \mu_s} = \frac{1}{2}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{1}{\lambda_h+\mu_h}$$

All terms are positive.

A.3.2 T^*

Claim: T^* **is decreasing in** $\lambda_h + \mu_h$**.**

$$\frac{\partial T^*}{\partial \lambda_h + \mu_h} = \frac{1}{\lambda_h + \mu_h(\lambda_s - (\lambda_h+\mu_h))} + \frac{1}{(\lambda_s - (\lambda_h+\mu_h))^2}\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)$$

Recall that the condition for $T^* \geq 0$ is

$$\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right) \leq 0$$

Furthermore, since $\lambda_h + \mu_h > \lambda_s$, the first term is negative. The derivative is the sum of two negative expressions and is negative.

Claim: T^* **is decreasing in** μ_s**.**

$$\frac{\partial T^*}{\partial \mu_s} = \frac{1}{(\lambda_s - (\lambda_h+\mu_h)(\mu_s+\lambda_h+\mu_h)}$$

$(\lambda_s - (\lambda_h+\mu_h))$ is negative and $(\mu_s+\lambda_h+\mu_h)$ is positive, so the sign of the derivative is negative.

Claim: T^* **is increasing in** λ_s**.**

$$\frac{\partial T^*}{\partial \lambda_s} = \frac{-1}{(\lambda_s - H)^2}\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{(\lambda_h+\mu_h)+\mu_s}{(\lambda_h+\mu_h)}\right)$$

This is the product of two negative terms, so it is positive.

A.3.3 $\overline{T}^*$

Claim: $\overline{T}^*$ **is decreasing in** $\lambda_h + \mu_h$**.**

$$\begin{aligned}\frac{\partial \overline{T}^*}{\partial \lambda_h + \mu_h} =& \frac{1}{(\lambda_s+\mu_s - (\lambda_h+\mu_h))^2}\left[\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right) + \mu_s\bar{t}\right] \\ &+ \frac{1}{(\mu_s+\lambda_s-(\lambda_h+\mu_h))(\lambda_h+\mu_h)}\end{aligned}$$

Because $\overline{T}^* \geq 0$, we must have $\left[\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)+\mu_s\bar{t}\right] < 0$, we have the first term is negative. By the assumption that $\lambda_h + \mu_h > \lambda_s + \mu_s$, the second term is negative. Therefore, this derivative is the sum of two negative terms, so it is negative.

Claim: **When $\bar{t}$ is sufficiently low, $\overline{T}^*$ is increasing in μ_s .**

$$\frac{\partial \overline{T}^*}{\partial \mu_s} = \frac{-1}{(\lambda_s+\mu_s-(\lambda_h+\mu_h))^2}\left[\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{(\lambda_h+\mu_h)+\mu_s}{(\lambda_h+\mu_h)}\right)+\mu_s\bar{t}\right] + \frac{1}{\lambda_s+\mu_s-(\lambda_h+\mu_h)}\left(\frac{1}{\mu_s+(\lambda_h+\mu_h)}+\bar{t}\right)$$

This is positive if

$$\bar{t} < T^* - \frac{\lambda_s+\mu_s-(\lambda_h+\mu_h)}{(\mu_s+(\lambda_h+\mu_h))(\lambda_s-(\lambda_h+\mu_h))}$$

Claim: $\overline{T}^*$ **is increasing in λ_s.**

$$\frac{\partial \overline{T}^*}{\partial \lambda_s} = \frac{-1}{(\lambda_s+\mu_s-H)^2}\left[\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{H+\mu_s}{H}\right)+\mu_s\bar{t}\right]$$

Because $\overline{T}^* \geq 0$, we must have $\left[\ln\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{H+\mu_s}{H}\right)+\mu_s\bar{t}\right] < 0$. Thus this is the product of two negative terms, so it must be positive.

B Proofs for welfare in the symmetric setting

B.1 Derivation of the probability of avoidable miscoordination

Recall that avoidable miscoordination arises when, conditional on one group being a hard type and one group being a soft type, the soft type makes the first commitment to its preferred alternative. This arises two ways: Firstly, the game progresses to T^* with no arrivals of λ_s, λ_h or μ_h. Secondly, the soft type is revealed at some time $\bar{t}$ and the game proceeds to $T^*(\bar{t})$ with no arrivals of λ_h or μ_h. In either case, μ_s arrives first after the relevant threshold is passed.

The probability of the first case arising is

$$e^{(-\mu_h-\lambda_h-\lambda_s)T^*}\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}$$

Plugging in the T^* derived earlier, this equals

$$\frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{-\lambda_s-\lambda_h-\mu_h}{\lambda_s-\lambda_h-\mu_h}}$$

This, however, does not account for possibility that $T^* = 0$. To include this possibility, the probability of avoidable miscoordination in the T^* case is

$$\begin{cases} \frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} & \text{if } T^* > 0 \\ \frac{\mu_s}{\mu_h+\lambda_h+\mu_s} & \text{if } T^* = 0 \end{cases} \tag{A17}$$

The probability of avoidable miscoordination in the $\overline{T}^*$ case is more complex, since $\overline{T}^*$ is not a fixed threshold, but a function of the time the soft type was leaked, $\bar{t}$. The probability of avoidable miscoordination in this case is given by

$$\int_{\tilde{t}=0}^{T^*} e^{(-\mu_h-\lambda_h-\lambda_s)\tilde{t}}(\lambda_s)e^{(-\mu_h-\lambda_h)(\overline{T}^*(\tilde{t})-\tilde{t})}\frac{\mu_s}{\mu_h+\lambda_h+\lambda_s}d\tilde{t}$$

Evaluating the integral yields, after simplification,

$$\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\frac{\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}{\mu_s(\mu_h+\lambda_h)-\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}\Bigg[\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} - \left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{(\mu_h+\lambda_h)}{\lambda_h+\mu_h-\lambda_s-\mu_s}}\Bigg] \tag{A18}$$

Therefore, the probability of avoidable miscoordination is, conditional on $T^* > 0$,

$$\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\Bigg[\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} + \frac{\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}{\mu_s(\mu_h+\lambda_h)-\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}\Bigg(\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} - \left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s-\mu_s}}\Bigg)\Bigg] \tag{A19}$$

and conditional on $T^* = 0$,

$$\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}\Bigg[1 + \frac{\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}{\mu_s(\mu_h+\lambda_h)-\lambda_s(\lambda_h+\mu_h-\lambda_s-\mu_s)}\Bigg(\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} - \left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s-\mu_s}}\Bigg)\Bigg] \tag{A20}$$

B.2 Proofs of comparative statics

Claim: **The probability of avoidable miscoordination is increasing in $(1-p)$, $\frac{u(A)-u(B)}{u(B)-u(SQ)}$.** Recall that increasing either of these terms decreased the duration of delay by an uninformed soft player. Since these terms only factor into the probability of avoidable miscoordination through the expressions for delay, and I have previously shown that the probability of avoidable miscoordination is decreasing in delay, the probability of avoidable miscoordination must be increasing in either of these factors.

Claim: **The probability of avoidable miscoordination is decreasing in λ_s.**
Fix $\lambda_0 > 0, \lambda' > \lambda_s$, and $\lambda_h > \lambda'$. Consider the equilibrium strategy played by soft types when λ_0 is the true rate of type revelation for soft types; denote by T^* and $\overline{T}^*$ the time thresholds that correspond to this equilibrium strategy as previously defined. We have already shown that these thresholds are increasing in the rate of type revelation for soft types. Therefore, we know that the equilibrium strategy played by soft types when λ' is the rate of information revelation involves time thresholds T', $\overline{T}'$ which are higher than T^*, $\overline{T}^*$, respectively.

Suppose we change the value of λ_s perceived by both players from λ_0 to λ' without changing the actual model value of λ_s. A hard type or a soft type who knows their opponent's type do not change their behavior. However, an uninformed soft type will choose to delay longer (T^* and $\overline{T}^*$ increase.) The interval of time $[T^*, T']$ provides an additional period during which types may be revealed (at the true rates). Therefore, this *directly* decreases the probability of avoidable negotiation failure by making it more likely that the hard type is revealed. It also *indirectly* decreases the probability of avoidable negotiation failure: suppose that during $[T^*, T']$, the λ_0-rate process arrives. This causes an uninformed soft type to postpone further to $\overline{T}' > \overline{T}^*$. During this additional interval of delay $[\overline{T}^*, \overline{T}']$, λ_h could arrive, which would render miscoordination impossible.

Now suppose we change the value of λ_s in the underlying model from λ_0 to λ', but players still believe that λ_0 is the true rate. This does not change the value of T^*, but makes it more likely that the soft type will be revealed before T^*. This furthermore makes it more likely that λ_h will arrive during $[T^*, \overline{T}^*]$. Therefore, this also indirectly decreases the probability of avoidable negotiation failure.

Now suppose we change both the perceived and true model value of λ_s from λ_0 to λ'. There a higher probability that the soft type could be revealed during $[0, T^*]$. Furthermore, there are histories after T^* during which the soft type would have started committing to his preferred alternative, but is now delaying until the new threshold T'. During $[T^*, T']$, there is also some likelihood that λ_h will arrive, which would rule out miscoordination, or that λ' will arrive, which will induce further delay until $\overline{T}'$, during which interval of delay λ_h could arrive and rule out miscoordination. All of these effects move in the direction of reducing the probability of avoidable miscoordination.

Claim: **In the T^* case, the probability of avoidable miscoordination is increasing in μ_s.**

Recall that the probability of avoidable miscoordination in the T^* case given in equation A17 was:

$$\begin{cases} \frac{\mu_s}{\mu_h+\lambda_h+\mu_s}\left(\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}} & \text{if } T^*>0 \\ \frac{\mu_s}{\mu_h+\lambda_h+\mu_s} & \text{if } T^*=0 \end{cases} \tag{A17}$$

Consider the second case. Differentiating with respect to μ_s yields $\frac{\mu_h+\lambda_h}{(\mu_h+\lambda_h+\mu_s)^2}$ which is clearly positive. Now consider the first case. The exponentiated term is also increasing in μ_s, since $\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}$ is increasing in μ_s and all of the other terms are positive and constant in μ. Therefore the probability of avoidable miscoordination is increasing in μ in the T^* case.

Proof of Remark 2.

The claim that when $\lambda_s=0$, the probability of avoidable miscoordination is increasing in μ_s is essentially already proven. When $\lambda_s=0$, the probability of entering the $\overline{T}^*$ case is 0 and so the proof follows from the result that in the T^* case, the probability of avoidable miscoordination is increasing in μ_s.

I next address the claim that when $\lambda_s=0$, the probability of avoidable miscoordination is decreasing in $\lambda_h+\mu_h$. Note first that $\frac{\mu_s}{\lambda_h+\mu_h+\mu_s}$ is decreasing in $\lambda_h+\mu_h$, that $\frac{1-p}{2p}\frac{u(A)-u(B)}{u(B)-u(SQ)}$ is constant in $\lambda_h+\mu_h$, that $\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}$ is increasing in $\lambda_h+\mu_h$, and $\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}$ is decreasing in $\lambda_h+\mu_h$. Therefore, as long as we can prove that

$$\left(\frac{\lambda_h+\mu_h+\mu_s}{\lambda_h+\mu_h}\right)^{\frac{\lambda_s+\lambda_h+\mu_h}{\lambda_h+\mu_h-\lambda_s}}$$

is decreasing in $\lambda_h+\mu_h$, then we are done. For ease of notation, I let $\Lambda\equiv\lambda_h+\mu_h$ for the rest of the proof. Differentiating the above expression with respect to Λ yields:

$$\left(\frac{\Lambda+\mu_s}{\Lambda}\right)^{\frac{\lambda_s+\Lambda}{\Lambda-\lambda_s}}\left(\frac{-2\lambda_s}{(\Lambda-\lambda_s)^2}\log\left(\frac{\Lambda+\mu_s}{\Lambda}\right)+\left(\frac{\lambda_s+\Lambda}{\Lambda-\lambda_s}\frac{-\mu_s}{\Lambda^2}\frac{\Lambda}{\Lambda+\mu_s}\right)\right)$$

All terms are positive except $\frac{-2\lambda_s}{(\Lambda-\lambda_s)^2}$ and $\frac{-\mu_s}{\Lambda^2}$. Therefore the sign of the derivative is negative.

C Proofs for equilibrium in the asymmetric setting

C.1 Derivation of the asymmetric equilibrium

I show the derivation of the best response for group a (it is symmetric for b). I first consider the case that both players are equally uninformed and no commitment have been observed by the threshold time. I retain the notation of $\Lambda^i \equiv \lambda_h^i + \mu_h^i$. Taking the opponent's optimal choice of delay T_b^* as given, T_a^* is either before or after T_b^* (the case of equality is addressed in the symmetric derivation). $T_a^*(T_b^*) < T_b^*(T_a^*)$ occurs when $T_b^* > K_a$, where K_a is the value in the domain that correspond to kinks in each player's best response. The case when best responses cross on the kinks is already done in the symmetric setting.

Case 1: $T_a^*(T_b^*)|T_b^* > K_a$. Then, $u(\text{commit to } A \text{ at } T_a^*) = (1-p_a)e^{-\lambda_s^b T_a^*}u_a(A)+p_a e^{-\Lambda^b T_a^*}u_a(SQ)$, and the continuation of waiting is the sum of the following cases:

1. b is a soft type, and a gets an arrival of μ_s^a in the interval $[T_a^*, T_b^*]$. A is implemented. This occurs with probability $(1-p_a)e^{-\lambda_b^s T_a^*} - (1-p_a)e^{(\mu_s^a-\lambda_s^b)T_a^* - \mu_s^a T_b^*}$
2. b is a soft type, and a doesn't get an arrival of μ_s^a in the interval $[T_a^*, T_b^*]$. Either A or B is implemented. The probability A is implemented is $(1-p_a)e^{(\mu_s^a-\lambda_s^b)T_a^* - \mu_s^a T_b^*}\frac{\mu_s^a}{\mu_s^a+\mu_s^b}$. The probability B is implemented is $(1-p_a)e^{(\mu_s^a-\lambda_s^b)T_a^* - \mu_s^a T_b^*}\frac{\mu_s^b}{\mu_s^a+\mu_s^b}$.
3. b is a hard type. Either SQ or B is implemented. The probability SQ is implemented is $(p_a)e^{-\Lambda^b T_a^*}\frac{\mu_s^a}{\Lambda^b+\mu_s^a}$. The probability B is implemented is $(p_a)e^{-\Lambda^b T_a^*}\frac{\Lambda^b}{\Lambda^b+\mu_s^a}$

Setting equal to a's utility of committing to A immediately, I obtain the threshold:

$$T_a^*(T_b^*)|T_b^* > K_a = \frac{1}{\lambda_s^b - \Lambda^b - \mu_s^a}\left(\ln\left[\frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)}\frac{\mu_s^b}{\mu_s^a+\mu_s^b}\frac{1-p_a}{p_a}\frac{\Lambda^b+\mu_s^a}{\Lambda^b}\right] - \mu_s^a T_b^*\right) \tag{A21}$$

Case 2: $T_a^*(T_b^*)|T_b^* < K_a$. Then, $u(\text{commit to } A \text{ at } T_a^*) = \left((1-p_a)e^{(\mu_s^b-\lambda_s^b)T_a^*+\mu_s^b T_b^*}\right)u_a(A) + (p_a e^{-\Lambda^b T_a^*})u_a(SQ)$, and the continuation value of waiting is the sum of the following cases:

1. b is a soft type. The first player to receive a commitment opportunity after T_a^* implements their preferred policy. The probability that A is implemented is $(1-p_a)e^{(\mu_s^b-\lambda_s^b)T_a^*+\mu_s^b T_b^*}\frac{\mu_s^a}{\mu_s^b+\mu_s^a}$. The probability that B is implemented is $(1-p_a)e^{(\mu_s^b-\lambda_s^b)T_a^*+\mu_s^b T_b^*}\frac{\mu_s^b}{\mu_s^b+\mu_s^a}$
2. b is a hard type. If a receives the first commitment opportunity, SQ stays in place. If b receives the first commitment, B is implemented. The probability that SQ is implemented is $p_a e^{-\Lambda^b T_a^*}\frac{\mu_s^a}{\Lambda^b+\mu_s^a}$. The probability that B is implemented is $p_a e^{-\Lambda^b T_a^*}\frac{\mu_s^a}{\Lambda^b+\mu_s^a}\frac{\Lambda^b}{\Lambda^b+\mu_s^a}$

Setting equal to a's utility of committing to A immediately, I obtain the threshold:

$$T_a^*(T_b^*)|T_b^* < K_a = \frac{1}{\lambda_s^b + \mu_s^b - \Lambda^b}\left(\ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right] + \mu_s^b T_b^* \right) \tag{A22}$$

Corresponding expressions for T_b^* can be derived symmetrically:

$$T_b^*(T_a^*)|T_a^* < K_b = \frac{1}{\lambda_s^a + \mu_s^a - \Lambda^a}\left(\ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] + \mu_s^a T_a^* \right) \tag{A23}$$

$$T_b^*(T_a^*)|T_a^* > K_b = \frac{1}{\lambda_s^a - \Lambda^a - \mu_s^b}\left(\ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] - \mu_s^b T_a^* \right) \tag{A24}$$

K_a and K_b are given by:

$$K_a := \frac{\mu_s^a + \mu_s^b}{\mu_s^a(\lambda_s^a - \Lambda^a - \mu_s^b) + \mu_s^b(\lambda_s^a + \mu_s^a - \Lambda^a)} \ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] \tag{A25}$$

$$K_b := \frac{\mu_s^b + \mu_s^a}{\mu_s^b(\lambda_s^b - \Lambda^b - \mu_s^a) + \mu_s^a(\lambda_s^b + \mu_s^b - \Lambda^b)} \ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right] \tag{A26}$$

Plugging each best correspondence function into the opponent's best correspondence function (for the same condition) yields closed-form expressions for equilibrium delay times:

$$T_a^*|T_a^* < T_b^* = \frac{\lambda_s^a + \mu_s^a - \Lambda^a}{(\lambda_s^a + \mu_s^a - \Lambda^a)(\lambda_s^b - \Lambda^b - \mu_s^a) + (\mu_s^a)^2} \ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right]$$
$$- \frac{\mu_s^a}{(\lambda_s^a + \mu_s^a - \Lambda^a)(\lambda_s^b - \Lambda^b - \mu_s^a) + (\mu_s^a)^2} \ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] \tag{A27}$$

$$T_b^*|T_a^* < T_b^* = \frac{\mu_s^a}{(\lambda_s^a + \mu_s^a - \Lambda^a)(\lambda_s^b - \Lambda^b - \mu_s^a) + (\mu_s^a)^2} \ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right]$$
$$+ \frac{\lambda_s^a + \mu_s^a - \Lambda^a}{(\lambda_s^a + \mu_s^a - \Lambda^a)(\lambda_s^b - \Lambda^b - \mu_s^a) + (\mu_s^a)^2} \ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] \tag{A28}$$

$$T_a^*|T_b^* < T_a^* = \frac{\mu_s^b}{(\lambda_s^a - \Lambda^a - \mu_s^b)(\lambda_s^b + \mu_s^b - \Lambda^b) + (\mu_s^b)^2} \ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right]$$
$$+ \frac{\lambda_s^a - \Lambda^a - \mu_s^b}{(\lambda_s^a - \Lambda^a - \mu_s^b)(\lambda_s^b + \mu_s^b - \Lambda^b) + (\mu_s^b)^2} \ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right] \tag{A29}$$

$$T_b^*|T_b^* < T_a^* = \frac{\lambda_s^b + \mu_s^b - \Lambda^b}{(\lambda_s^a - \Lambda^a - \mu_s^b)(\lambda_s^b + \mu_s^b - \Lambda^b) + (\mu_s^b)^2} \ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1 - p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right]$$
$$- \frac{\mu_s^b}{(\lambda_s^a - \Lambda^a - \mu_s^b)(\lambda_s^b + \mu_s^b - \Lambda^b) + (\mu_s^b)^2} \ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1 - p_a}{p_a} \frac{\Lambda^b + \mu_s^a}{\Lambda^b}\right] \tag{A30}$$

C.2 Proof of Proposition 4

Claim: $T_a^*(T_b^*)$ and $T_b^*(T_a^*)$ can only intersect once.

The 45 degree line $T_a^* = T_b^*$ divides the best response space into two regions: $T_a < T_b$ and $T_a > T_b$. The following lemma describes bounds on the behavior of both best response correspondences in each region, which will then be used to prove the statement of uniqueness.

Lemma A1. *$T_i^*(T_j^*)$ kinks when $T_i = T_j$ and has linear subfunctions. When $T_i < T_j$, the slope of $T_i^*(T_j^*) \in (0, 1)$. When $T_i > T_j$, the slope of $T_i^*(T_j^*) \in (-\infty, 0)$.*

Proof. Note that for either player i, the functions that describe i's best response $T_i(T_j)$ conditional upon $T_i > T_j$ and conditional upon $T_i < T_j$ are both linear in T_j.
The slope of i's best response $T_i(T_j)$ conditional upon $T_j < T_i$ is

$$\frac{\mu_s^i}{\lambda_s^j - H^j - \mu_s^i} = \frac{-\mu_s^i}{H^j - \lambda_s^j + \mu_s^i}$$

Since $H^j > \lambda_s^j$ and $\mu_s^i > 0$, this must be in $(-\infty, 0)$.

The slope of i's best response $T_i^*(T_j)$ conditional upon $T_i^* < T_j$ is

$$\frac{-\mu_s^j}{\lambda_s^j + \mu_s^j - H^j} = \frac{\mu_s^j}{H^j - (\lambda_s^j + \mu_s^j)}$$

Again, since $H^j > \lambda_s^j + \mu_s^j$, this is bounded in (0, 1).

Because both intervals are open, T_i cannot have zero slope in either region, so its slope must change at $T_i = T_j$. At the identity, the value of both functions equals $T_i(K_j)$, thus there must be a kink, rather than a discontinuity. This proves the Lemma. □

Case 1. Suppose players' best response functions intersect on the 45 degree line. This must mean that the functions cross exactly on each of their kinks. Then, we can only have multiple crossings if $T_i^*(T_j^*)$ has the same slope as $T_j^*(T_i^*)$ in one or both of the regions. However, when $T_i^* > T_j^*$, the slope of $T_i^*(T_j^*) \in (-\infty, 0)$. The slope of $T_j^*(T_i^*)$ is $\frac{\mu_i}{H^j - (\lambda_i + \mu_i)}$. Inverting to represent T_i as a function of T_j, we have $\frac{H^j - (\lambda_i + \mu_i)}{\mu_i}$ which is between $(1, \infty)$. This interval is disjoint from $(-\infty, 0)$, so there is no intersection. Similar reasoning applies when $T_i^* < T_j^*$: The slope of $T_i^*(T_j^*)$ is constrained between $(0, 1)$. The slope of $T_j^*(T_i^*)$, inverted to represent T_i as a function of T_j, is constrained between $(-\infty, -1)$. Conditional upon intersection at $T_i = T_j$, there cannot be an intersection in this region.

Case 2. Suppose that the players' best responses kink in different places. Assume WLOG $K_i < K_j$. Note that the order of the kinks implies that $T_i^*(K_j) < T_j^*(K_j) = K_j$. We will show that the functions must cross once in the region where $T_i^* < T_j^*$. In this region, the slope of $T_i^*(T_j^*)$ is in $(0, 1)$, while the slope of $T_j^*(T_i^*)$, inverted to be a function of T_j^*, is negative. Given that $T_i^*(T_j^*)$ is below $T_j^*(T_i^*)$ at K_j, $T_i^*(T_j^*)$ is an increasing function (that stays below the 45-degree line), and the inverted $T_j^*(T_i^*)$ is a decreasing function, they must intersect.

To see why they cannot cross in the region where $T_j^* > T_i^*$, note that $T_i^*(K_j) < T_j^*(K_j) = K_j$. In this region, $T_j^*(T_i^*)$ (as a function of T_j^*) has an positive while $T_i^*(T_j^*)$ has negative slope. Furthermore, $T_i^*(T_j^*)$ must intersect with $K_i < K_j$. Therefore they must diverge in this region.

C.3 Supplemental result: $\overline{T}^*$ in the asymmetric case

Proposition A1 (Commitment delay for leaked soft types)**.** *Define a relevant history as one where neither group has committed to an alternative and neither group's type has been leaked. Consider the continuation game at a relevant history. For each player $i = a, b$, there exists a unique threshold time T_i^* such that if both groups' types are unknown, a soft group will commit* *to its preferred alternative if and only if $t > T_i^*$.*

Now consider a sub-history of h where i's type is revealed at time $\bar{t}_i < \min\{T_a^*, T_b^*\}$. Then, there exists a unique $\overline{T}_i^*(\bar{t}_i) > \min\{T_a^*, T_b^*\}$ such that a soft group will commit to its preferred alternative iff $t > \overline{T}_i^*(\bar{t}_i)$.

Proof: Consider $\overline{T}_a^*$. The proposition claims that *timing* of leaks has an additional restriction: $\bar{t}_i < \min\{T_a^*, T_b^*\}$. That is, a soft type being leaked matters only *before the earlier of the two thresholds.* To see why, consider the problem faced by a soft type of group a. First suppose $T_b^* < T_a^*$ and a's type is revealed during the interval $[T_b^*, T_a^*]$. Group b is already trying to commit by the time a's type is leaked. Therefore, a's type being leaked will not affect a's learning about b's type. Now suppose that $T_a^* < T_b^*$ and a's type is leaked during the interval $[T_a^*, T_b^*]$. At this point, a has already passed the threshold where they are sufficiently confident that b is a soft type, and therefore is already trying to commit. Having a's own type leaked does not affect inferences that a has already made about b's type, and hence does not affect a's behavior. Hence, $\overline{T}_a^*$ is a function of $\bar{t}_a$ and not of T_b^*, as a's type being leaked induces b to start committing as soon as possible, making T_b^* irrelevant.

Group a's expected utility of committing to A at $\overline{T}_a^*$ is

$$\left[(1-p_a)e^{-\lambda_s^b \overline{T}_a^* - \mu_s^b(\overline{T}_a^* - \bar{t}_a)}\right] u_a(A) + \left[p_a e^{-\Lambda^b \overline{T}_a^*}\right] u_a(SQ)$$

The continuation value of waiting is:

$$(1-p_a)e^{-\lambda_s^b \overline{T}_a^* - \mu_s^b(\overline{T}_a^* - \bar{t}_a)}\left[\frac{\mu_s^a}{\mu_s^a + \mu_s^b} u_a(A) + \frac{\mu_s^b}{\mu_s^a + \mu_s^b} u_a(B)\right] +$$
$$p_a e^{-\Lambda^b \overline{T}_a^*}\left[\frac{\mu_s^a}{\mu_s^a + \Lambda^b} u_a(SQ) + \frac{\Lambda^b}{\mu_s^a + \Lambda^b} u_a(B)\right]$$

Setting these equal and rearranging, I obtain

$$\overline{T}_a^* = \frac{1}{\lambda_s^b + \mu_s^b - H^b}\left(\ln\left[\frac{u_a(A) - u_a(B)}{u_a(B) - u_a(SQ)} \frac{\mu_s^b}{\mu_s^a + \mu_s^b} \frac{1-p_a}{p_a} \frac{H^b + \mu_s^a}{H^b}\right] + \mu_s^b \bar{t}_a\right) \tag{A31}$$

A symmetric derivation shows that

$$\overline{T}_b^* = \frac{1}{\lambda_s^a + \mu_s^a - \Lambda^a}\left(\ln\left[\frac{u_b(B) - u_b(A)}{u_b(A) - u_b(SQ)} \frac{\mu_s^a}{\mu_s^a + \mu_s^b} \frac{1-p_b}{p_b} \frac{\Lambda^a + \mu_s^b}{\Lambda^a}\right] + \mu_s^a \bar{t}_b\right) \tag{A32}$$

Thus, $\overline{T}_i^*$ is given by:

$$\overline{T}_i^* = \frac{1}{\lambda_s^j + \mu_s^j - (\lambda_h^j + \mu_h^j)}\left[\ln\left(RD_i \frac{\mu_s^j}{\mu_s^i + \mu_s^j} \frac{1-p_i}{p_i} \frac{(\lambda_h^j + \mu_h^j) + \mu_s^i}{(\lambda_h^j + \mu_h^j)}\right) + \mu_s^j \bar{t}_i\right] \tag{A33}$$

D Supplemental figures

D.1 Welfare

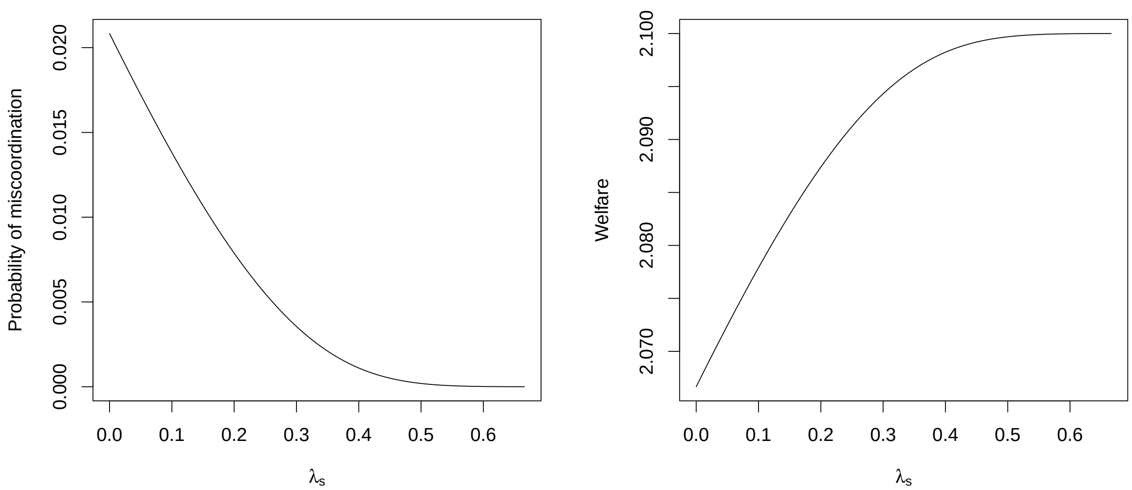

(a) Miscoordination **(b)** Welfare

Figure A1. Effects of changing λ_s on miscoordination and welfare.
Parameter values: $\mu_s = \frac{1}{3}, \lambda_h + \mu_h = 1, \frac{1-p}{p} = \frac{1}{4}, \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{1}{2}$

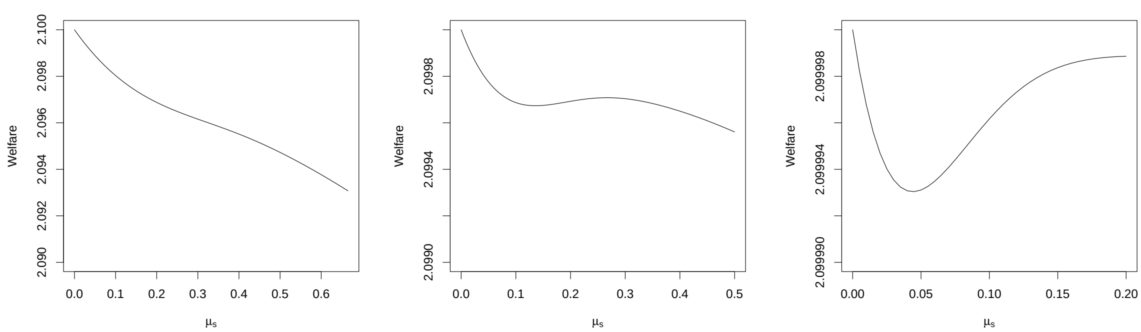

(a) $\lambda_s = \frac{1}{3}$ **(b)** $\lambda_s = \frac{1}{2}$ **(c)** $\lambda_s = \frac{2}{3}$

Figure A2. Effects of changing μ_s on welfare, as conditioned by λ_s. Panel (a) has parameters identical to those used in Figure 4. As I progressively increase λ_s in panels (b) and (c), the likelihood that one-sided asymmetry will be triggered increases, and comes to dominate the aggregate effect. (Note that increasing λ_s curtails the range of possible values for μ_s.)
Parameter values: $\lambda_h + \mu_h = 1, \frac{1-p}{p} = \frac{1}{4}, \frac{u_a(A)-u_a(B)}{u_a(B)-u_a(SQ)} = \frac{1}{2}$

D.2 Asymmetric setting

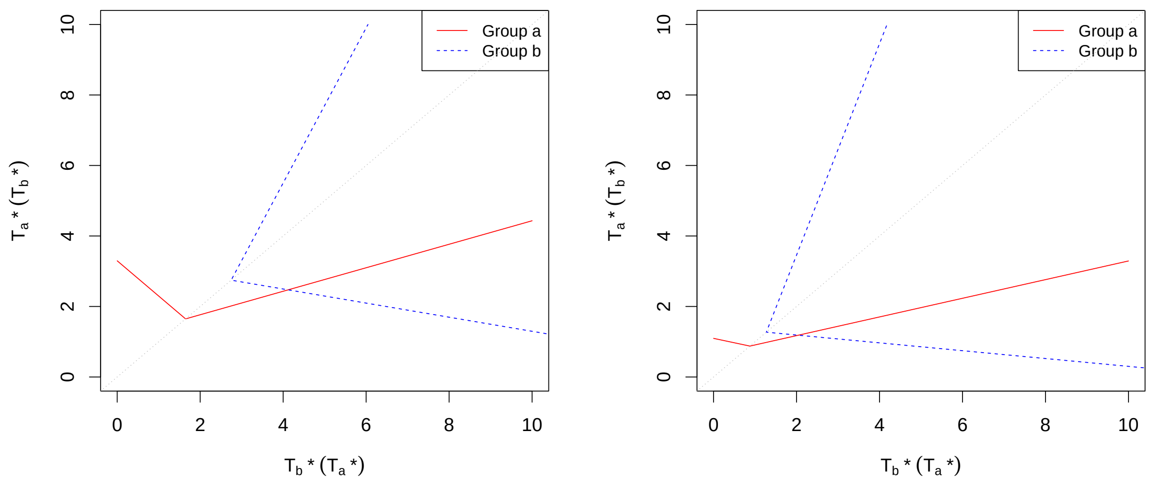

(a) $\lambda_s^a = 0.6$ **(b)** $\mu_s^a = 0.6$

Figure A3. Effects of higher λ_s^b and μ_s^b on equilibrium delay. Other parameter values are as in Figure 2.

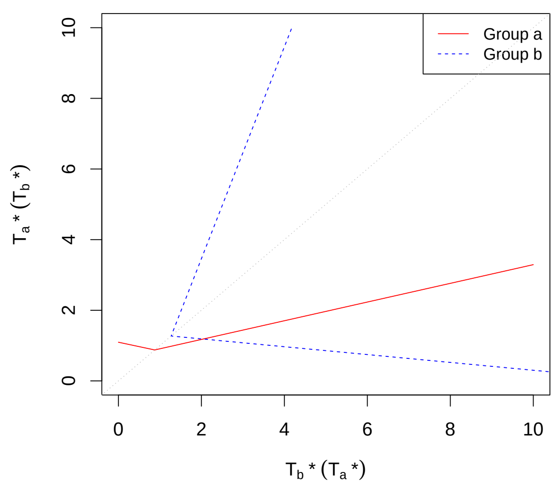

Figure A4. Effect of higher $\lambda_h^b + \mu_h^b$ on best responses. $\lambda_h^b + \mu_h^b = 2$. Other parameter values are as in Figure 2.