Preemptive Offers versus Counter Offers with Investments in Fallback Positions*

Kemal Kıvanç Aköz† Nejat Anbarcı‡ Kang Rong§

March 18, 2018

Abstract

We consider a setup where at any point Player 1 can make Player 2 an offer, who can have one of the two possible levels of an initial fallback position/payoff, high or low, and can also make an investment to improve her fallback position without improving the joint surplus available to players. The total of 2's initial fallback position and 2's investment is her final fallback position. In the complete-information case, the final fallback position is fully observable by 1, and it turns out that the outcome is robust to whether 1 can make at most one offer or multiple offers instead. 1 can either wait for 2 to make her optimal investment and then match any outside offer she receives by a counter offer or instead make a preemptive offer before 2 invests. A particular preemptive offer turns out to be fully efficient while a counter offer is never efficient. The payoff difference between the two surplus levels is fully channeled to 1. In the no-information case, nothing is observable by 1. In the noisy-information case, 1 can receive two separate informative but noisy signals. One signal is about 2's initial bargaining position, and the other about 2's investment level. The outcomes in the no- and noisy-information cases are not robust to whether 1 can make at most one offer or multiple offers. If 1 can make at most one offer, we have the following results in these cases. When 1's prior belief about 2's high type is high, 1 makes

*We would like to thank audiences at the University of Sussex, NYU at Abu Dhabi, Sabanci University, Bilkent University, University of Crete, Middle-Eastern Technical University, Ozyegin University, Barcelona JOCS seminar, and the IX. Winter Workshop in Economics (at Koc University).

†New York University Abu Dhabi (e-mail: kka227@nyu.edu).

‡Deakin University (e-mail: nejat.anbarci@deakin.edu.au)

§Shanghai University of Finance and Economics (e-mail: rong.kang@mail.shufe.edu.cn)

a preemptive offer that 2's both types accept. It is efficient. When 1's prior belief about 2's high type is not very decisive and the cost of Player 2's optimal investment is low, 1 makes a counter offer, which is not efficient. In general, 1 makes a signal-contingent preemptive offer after observing the signal about 2's type. It turns out that Player 1 chooses a counter offer over a signal-contingent preemptive offer as the cost of Player 2's optimal investment goes down. When 1 is allowed to make multiple offers, 1 typically combines the low preemptive offer with the high counter offer. In the no-information (noisy-information) case, for a relatively large range of relatively low prior beliefs (low signals) that he is facing the high type of 2, Player 1 makes the low preemptive offer which only the low type of 2 accepts; if this offer is rejected, then he makes the appropriate counter offer to 2 after she receives offer from the market. If, however, 1 has (receives) a sufficiently high prior belief (high signal) and the cost of investment is high, to guarantee an agreement without any investment, 1 makes the high preemptive offer which both types of 2 accept. We also find that, regardless of whether 1 can make at most one offer or multiple offers, the noisy-information case will always be less conducive to a counter offer.

JEL: C72, C78, D86

Keywords: Investing in fallback positions, counter offer, preemptive offers, insurance and signal-manipulation incentives, efficiency, at most one offer vs. multiple offers.

1 Introduction

1.1 Motivation

In the modern property rights approach to the theory of the firm, developed by Grossman and Hart (1986) and Hart and Moore (1990), the hold-up problem is a central ingredient. In a nutshell, a party can make a non-contractible observable and irreversible investment, where the investment increases the surplus that can be generated within a given relationship more than it increases the party's fallback position/payoff. Since the investing party does not have all the bargaining power ex post, she cannot reap the full returns of her investment, and thus in general there is an under-investment problem. In this literature, it is a crucial assumption that investments are partially relationship-specific and partially person-specific. Consequently, ownership arrangements can only affect what a party can get outside of the relationship.

Another standard assumption of this literature is that there is symmetric information between the parties. As such, under symmetric information, bargaining will always be successful with an efficient agreement. Nevertheless, the parties' fallback payoffs are relevant, because they influence the agreement payoffs as investment incentives depend on these fallback payoffs. Thus, institutions, i.e., governance structures in terms of ownership arrangements, matter because of the symmetric information. The symmetric information assumption, however, is deemed "deeply problematic" by Williamson (2002, p.188), as a party may have better information than the other party about her fallback position.

In this paper, we reverse the crucial assumptions of the property rights literature. Specifically, we assume that (i) investments, which are still strategic and endogenous, are *completely person-specific* (i.e., not relation-specific at all) in that their effect on the joint surplus is zero, thus entirely falling on the parties' fallback positions, and (ii) information is typically asymmetric in that at least one party knows her fallback payoff more precisely than the other party. Our setup is then becomes conducive to analyze new topics like preemptive retention offers vs. ex-post counter offers - as our lead example below will illustrate[1] - as well as shedding a new light to age-old issues such as arms races (or defense expenditures) in a conflict, out-of-court settlements in costly litigation, lobbying, and opportunistic behavior in mergers and acquisitions, etc.

Arming, legal costs, on-the-job search and various other endogenous influence activities can give rise to endogenous threat payoffs and thus increase one party's fallback position. Nevertheless, since these activities are costly, they also deduct from what is obtainable for both parties. As a result of the reversal of the crucial assumptions of the property rights literature, in our setup there is a reversal of the main outcomes of that literature: consequently, the outcome is inefficient and, in addition, there is an over-investment instead of under-investment by parties. Moreover, the inefficiency is not only due to the over-investment which is a totally person-specific sunk cost, but it is also due to the loss of surplus brought up by the disagreement probability, which in turn arises as a result of asymmetric-information bargaining between parties.

Consider a faculty member at a university who is 'looking around' by investing in her fallback position, vis-à-vis her department head or dean - shortly the administrator -, by incurring an irreversible cost. Her endogenous investment can, for instance, be

[1]There is no academic literature on this topic in any discipline, while there are many non-academic or informal discussions of the topic on the internet and at academic institutions (see below).

in the form of giving seminars at other universities - some of them being explicit job talks - which may require preparation and travel time, among other costs. Clearly these activities, while not increasing the surplus between the administrator and faculty member, can nevertheless improve the level of an offer that the faculty member may generate from outside of her university (i.e., can improve her fallback position), which in turn can force the administrator to make a preemptive retention offer to the faculty member before investing in these activities or to wait and make a matching counter offer ex post, in case the faculty member manages to generate an outside offer after her investment activities.[2]

In terms of what the administrator knows about the faculty member's fallback position before and after her investment activities, one possibility is that the administrator is totally 'on top of things', which will be the very stylized complete-information case: That is, in the first place, he is able to observe how the market perceives the faculty member as a candidate even in the absence of her activities to invest to her fallback position - i.e., observe the faculty member's initial fallback position judged at least by head-hunters or pro-active hiring/recruitment committees of other departments in the profession. In addition, he also knows how much the faculty member's investment activities - i.e., the quantity and quality of her campus visits, whether it is simply a seminar presentation or a job talk, and in case it is a seminar presentation whether it can lead to an opportunity hiring in those particular departments, and in case it is a job talk the quality of the applicant field in those particular departments - and how much they can add to her initial fallback position (as perceived by outsiders) and how much these activities would cost her.

The other extreme (and also somewhat stylized) case would be where the administrator neither knows the faculty member's initial fallback positions - e.g., the administrator and the department may lack expertise in the faculty member's specific field and thus cannot assess her initial fallback position - nor her investment activities - e.g., the administrator and the department may not have a good network to hear about seminar and campus visit activities. The administrator will only know how much these activities would cost her. This will be the case of no - or hidden -

[2]Many leading universities now have explicit preemptive retention policies. The ones adopted by UCSD, University of Washington, and University of Oregon are provided in the following documents: http://soeadm.ucsd.edu/ppi/academic_personnel/files/docs/Retention_Guidelines_-_Criteria_for_Analysis_10-10-2016.pdf.; http://ap.washington.edu/ahr/policies/compensation/salary-adjustments/retention-salary-adjustments/ ; https://academicaffairs.uoregon.edu/retention-salary-adjustments

information.

A perhaps more realistic setup would be between these two extreme benchmark cases: The administrator neither knows the faculty member's initial fallback position exactly nor can observe the faculty member's endogenous investment activities exactly, but can obtain informative signals about both.

One can easily find other examples of our basic setup.[3] Further, our model can also fit the situations of conflict easily, especially if one interprets disagreement outcome as conflict outcome and fallback position as military capacity.[4] As such, we contribute to the literature on conflict by providing a model that endogenizes bargaining power as an outcome of strategic investment in fallback position and analyzes the impact of various information structures on key outcomes such as military capacity and likelihood of conflict (see Ramsay, 2017, for a recent literature survey on the role of uncertainty in conflicts). Gennaioli and Voth (2015) also provide a model on strategic investment on military capacity in their model on the trade-off between building fiscal and military capacity. However, they treat disagreement as an exogenous random event. Baliga and Sjöström (2013) offer a model that is closer to ours. They analyze investment for the military capacity that reduces the total surplus but increases the expected payoff of the disagreement outcome.

The closest work to ours is Anbarci, Skaperdas and Syropoulos (2002) who considered economic environments where agents make costly and irreversible endogenous, strategic investments - in "guns" - that are publicly observable and may enhance their respective threat payoffs without enhancing the surplus at all, where both parties also jointly adhere to a bargaining solution, i.e., a "norm", given their utility possibilities

[3]E.g., an office employee may start a costly on-the-job search and her administrator (or principal) can make a preemptive retention offer instead of starting an abrupt costly search again. In a more cooperative vein, suppose it is well known that a large company, intending to acquire a smaller (regional) company, is assessing the value of the small company, which in turn may feel tempted to invest in itself to increase its market appeal to other potential large companies/suitors (where the potential suitors may also consider acquiring it) in case the large company's assessment of it does not come out to its liking.

[4]Investment in military capacity determines the expected payoff from conflict. In line with this interpretation, an example from international relations involves investments in weapons or in weapon technology by a relatively small and closed country (e.g., Iran, North Korea), whose initial level of weapons or weapon technology may not exactly be observed by outsiders. A large world power, such as the U.S., may find it in its interest to reach a deal with that country preemptively instead of facing a disagreement or impasse which may be conducive to a potentially destructive conflict in the future. Another example could be bargaining between a state and an insurgent group over some of the resources in a territory. Our work sheds light on how strategic investment by the insurgent group on its military capacity changes with what the state knows about the insurgent group's investment.

set. In this complete-information setting, it then becomes possible to rank different bargaining solutions in terms of their efficiency.[5]

Nevertheless, Anbarci et al. (2002) clearly used a hybrid setup where the outcome is determined via a bargaining solution, - i.e., non-strategically -, upon the strategic decision of the parties regarding their investments in their threat payoffs, envisaging which bargaining solution is to decide on their outcome.

In this paper, we go much further than that. We consider a fully strategic setup which may or may not involve incomplete information about one party's initial fallback position and her strategic investment. One party may make endogenous and costly investments, and this party coincides with the one whose initial fallback position is not common knowledge, unless there is complete information about both parties' initial fallback positions. We first consider a setting where the other, uninformed, party makes an offer. With complete information, one party could be designated as the one who invests in her fallback position and the other party then is supposed to make an offer.

1.2 Brief Summary of Our Model and Results

In our analysis, there are two risk-neutral players, Players 1 and 2, who bargain over a fixed-size surplus, which is normalized to one. Player 2 can have one of the two possible levels of initial fallback position, namely high and low, and can make an investment to enhance her initial fallback position further so that the total of her initial fallback position and her investment will constitute her final fallback position. We assume that Player 2's final fallback payoff maps to the best offer from the market.

We consider two related setups. The first one assumes that Player 1 (i.e., the administrator) can make at most one offer to Player 2. But there is no restriction on when Player 1 can make an offer to Player 2. It can take place at any point of time, i.e., before or after receiving complete-, no- or noisy-information about Player 2's initial fallback position or before or after Player 2 makes an investment in her fallback position. If Player 2 accepts it, it is enforced; otherwise, they both receive

[5] Anbarci et al. (2002) compared bargaining solutions within a class - such as the Kalai-Smorodinsky solution, the Egalitarian solution and the Equal Sacrifice solution -, in which the influence of the threat point on the bargaining outcome varies across solutions. Under symmetry, they found that the solution in which the threat point is least influential, i.e., the Equal Sacrifice solution, Pareto-dominates the other solutions. Since the Equal Sacrifice solution puts the least weight on the threat point, they conclude that "norms" against threats can mitigate some of the costs of conflict and therefore have efficiency-enhancing effects.

their fallback payoffs. Thus, in a general sense, this is an ultimatum game with fallback payoffs where the final fallback payoff of Player 2 is determined by her investment.[6]

The motivation for this setup, which restricts the number of offers made by Player 1 to at most one offer, is as follows. In real life administrators typically refrain from making further offers after an offer they have made is already rejected by the faculty member. It could be partially because the other faculty members in the department may not want someone back in their department any more after that faculty member refuses the offer by their dean or department head (and thus by their faculty or department). Given that, the dean or department head may not want to make the department look too weak vis-à-vis that particular faculty member either via making a further offer to her. It could also be related to the faculty member's perceptions. She may not appreciate a low-balling offer even if it is later followed by a better one. She may feel that the faculty or department does not value her sufficiently highly; she may think that otherwise she would have received the higher offer right from the outset. So, she may reject the offer from the faculty or department and accept the outside offer even if the offer from the department matches the outside offer. In addition, for the faculty member there may be additional long-run costs of declining an outside offer since it may generate a long-lasting belief about her that she tends to go to the job market in order to get a counter offer from her university, i.e., that she is not serious about leaving her department, causing unnecessary cost to other universities that would be interested in her. Our second setup does not have any restrictions on the number of offer Player 1 can make to Player 2 even if she has declined any of them in the past.

Within our first setup with at most one offer, consider the complete-information case where Player 1 waits for Player 2 to make her optimal level of investment and receive offers from the market and then makes her a counter offer to match the best offer she has received from the market, assuming that Player 2's best offer does not exhaust the entire pie. In this case, given a tie-breaker assumption, Player 2 will accept this take-it-or-leave counter offer over the best offer from the market and Player 1 will end up receiving the entire expected net surplus (i.e., the surplus net of Player 2's final fallback payoff which maps to the best offer she receives). Thus, Player 1's counter

[6] One may think that our setup resembles that of job market signaling models to some extent. Note that there is one major difference: high and low types of Player 2 would imply a different size of surplus size in job market signaling in favor of the high type whereas the gross surplus sizes are the same for both types of Player 2 in our setup and the net surplus size with the high type is smaller than that with the low type.

offer yields zero net surplus to Player 2, where the expected size of the net surplus is 1 minus Player 2's final fallback payoff.

Trivially, the net surplus could become size 1 easily, if Player 1 ex-ante offered the entire net surplus in the form of a preemptive offer to Player 2 who would then accept it and consequently refrain from making any investment. Thus, this extremely generous preemptive offer by Player 1 can reduce Player 2's optimal investment level to zero, but it does not benefit Player 1 at all. Hence, Player 1 can search for other preemptive offers that he can make to Player 2 ex ante that would still alleviate - or even get rid of altogether - Player 2's investment while also benefiting Player 1 himself.

We first find that there is a particular preemptive offer that is the optimal offer for Player 1, which also turns out to be efficient while any other offer, including the counter offer, is not. Player 2 accepts this preemptive offer and consequently never gets to make any investment. With this particular offer, Player 2 receives the difference between the additional benefit of her investment minus the cost of her investment, which will term the 'investment surplus'. The payoff difference between the two surplus levels (one with the preemptive offer and the other with the counter offer), however, is fully channeled to Player 1. Later, it turns out that the complete-information case is robust to whether Player 1 can make at most one offer or multiple offers.

In the no-information case, nothing is known to Player 1, except for Player 2's investment cost function and a prior belief about the type of Player 2. When Player 1's prior belief about the high type of Player 2 is high, Player 1 chooses to make a preemptive offer that both types of Player 2 accept. The outcome then is efficient. With this offer, relative to the complete-information outcome, the high type of Player 2's payoff stays put, the low type of Player 2 gains, and Player 1 loses. When Player 1's prior belief about the high type of Player 2 is not very decisive and the cost of Player 2's optimal investment is low, however, instead of a preemptive offer Player 1 chooses to make a counter offer, which is not efficient. The intuition is as follows. Note that the advantage of making a counter offer is that Player 1 can precisely target different types of Player 2 (i.e., Player 1 can make different offers to different types of Player 2). Obviously, this advantage - of making a counter offer over a preemptive offer - is stronger if Player 1's belief about Player 2's type is not very decisive. On the other hand, the advantage of making a preemptive offer is that when Player 1 makes a preemptive offer, Player 2 will not made any investment yet and thus would be willing to accept lower offers in order to avoid the cost of investment. Obviously, this advantage - of making a preemptive offer over a counter offer - is stronger if the cost of investment is high. So, for Player 1, making a preemptive offer is more attractive than making a counter offer only if Player 1's belief about Player 2 is not very decisive and the cost of Player 2's (optimal) investment is low.

In the noisy-information case, Player 1 can receive two separate noisy but informative signals. One signal is about Player 2's initial fallback position. The other signal is about Player 2's investment level. Note that in this case Player 2 would make her investment not only to optimize her final fallback position, but also to 'manipulate' the signal that Player 1 receives and secure a decent outside offer from the market by investing on top of the initial bargaining in case Player 1 makes a very low offer; the latter is the insurance incentive. (The signal-manipulation channel will be especially appealing to the low type of Player 2 while the insurance channel will be especially appealing to the high type of Player 2.)

Note that in the noisy-information case, there are two more possibilities that Player 1 can choose from: signal-contingent preemptive offers and double-signal-contingent preemptive offers, as two separate signals - one signal about the type of Player 2 and, if any, another one about the investment level of Player 2 - could be observed by Player 1. We show that double-signal-contingent preemptive offers are never optimal for Player 1. We also show that signal-contingent preemptive offers dominate signal-noncontingent preemptive offers. Thus, it boils down to two types of viable offers for Player 2: signal-contingent preemptive offers vs. counter offers. It turns out that Player 1 chooses a counter offer over the signal-contingent preemptive offer as the cost of Player 2's optimal investment goes down.

When Player 1 is allowed to make multiple offers, in the no- and noisy information cases, Player 1 typically prefers to combine the low preemptive offer with the high counter offer. For a relatively large range of relatively low prior beliefs, he makes the low preemptive offer which only the low type of Player 2 accepts; if this offer is rejected, then he makes the appropriate counter offer to Player 2. after she receives offers from the market. If, however, Player 1 has sufficiently high priors that he is facing the high type of Player 2 and the cost of investment is high, Player 1 chooses to make the high preemptive offer to guarantee an agreement without any investment. In the noisy-information case, the only difference is that Player 1 makes his decisions contingent on the type signal he receives instead of on his priors. Note, however, that Player 1's posterior belief after observing the signal τ is more precise than his prior belief. Therefore, regardless of whether Player 1 can make at most one offer or multiple offers, in the noisy information case it is harder for Player 1 to remain too uncertain about the type of Player 2 to wait for the counter-offer opportunity; consequently, the noisy-information case will always be less conducive to a counter offer.

2 Other Related Literature

In this section we will elaborate on work that has not been mentioned in the Introduction. Fisher and Ury (1981), in their classic and best-selling book on negotiation, contended that parties in any bargaining would be wise to invest resources in enhancing their "best alternative to a negotiated agreement" (BATNA) which is their fallback option in case they fail to reach an agreement. This recommendation is supported by ample robust empirical evidence in negotiation and business administration literatures, showing that the more attractive is a party's best alternative to a negotiated agreement, the better is her fallback position and bargaining power (see Pinkley, Neale, and Bennett, 1994, and Pinkley, 1995, for instance).

Further, Mahotna and Gino (2011) have recently illustrated in a experiment how attempts aimed at enhancing one's fallback position can allow parties to obtain gains in their negotiations, even after controlling for the leverage provided by the outside options.[7] Their results demonstrate that previously sunk investments in generating an outside option lead to a magnified sense of entitlement, even when the outside option has already been foregone.

Very recently, Morita and Servatka (2013) have noted that investments by parties in their fallback positions may also be a source of ex-post opportunistic behavior in bilateral trade relationships; parties may exert effort to search for alternative business partners even if it does not add trade value. Morita and Servatka (2013) also noted that such investments might negatively affect the parties' other-regarding preferences if the investment is viewed as opportunistic. They experimentally investigated a bilateral trade relationship in which standard theory assuming self-regarding preferences predicts that the seller will be better off by investing in the outside option to improve his fallback position. They, however, found overall support for the link between other-regarding behavior and opportunism.

[7]There is a difference between outside options and disagreement payoffs. Outside options must be chosen in lieu of bargaining, instead of being available afterward in the event bargaining fails. Note that our setup involves disagreement payoffs rather than outside options.

Bargaining involving obstinate/commitment types also relates to our setup. Suppose that there is incomplete information about the type of Player 2. In particular, Myerson (1991) showed that if Player 2 is potentially a strategically inflexible "commitment" type that insists on portion θ_2 of the bargaining surplus, and Player 1 is a fully rational normal type with certainty, then Player 2 obtains θ_2 and Player 1 receives $1-\theta_2$ in any perfect equilibrium, even if the probability that Player 2 is a commitment type is arbitrarily small; this happens because rational normal types behave in a manipulative way, mimicking the commitment types. Note that, nevertheless especially the rational type of Player 2 incurs a cost by rejecting Player 1's offer in that she misses beneficial bargains that would yield her below θ_2. In addition to incomplete information about the type of Player 2, suppose that there is also incomplete information about the type of Player 1. In particular, as Abreu and Gul(2000) showed, if both players are potentially commitment types that demand θ_1 and θ_2, then a war of attrition ensues, and the unique equilibrium payoff profile is inefficient with the "weak" agent (Player i) receiving $1-\theta_j$ and the "strong" agent receiving strictly less than θ_j.

More recently, in Atakan and Ekmekci (2014), even if the frequency of behavioral types - which is determined in equilibrium - is negligible, they affect the terms of trade and efficiency. To be more specific, the magnitude of inefficiency is determined by the demands of the commitment types and, interestingly, is independent of their frequency. Thus, access to the market exacerbates bargaining inefficiencies, even at the frictionless limit of complete rationality.

3 The Model: Preliminaries

There are two risk-neutral Players 1 and 2, who bargain over a fixed size of a surplus (or pie), which is normalized to one. (We will continue referring to Player 2 as "she" and Player 1 as "he" in what follows.) Player 2 is one of the two types $\{L, H\}$; i.e., she either has a low initial fallback payoff $d_L = 0$ or a high one, $d_H \in (0, 1)$. Player 2's type is privately known by her, unless it is the complete-information case in which case all types are publicly known. Player 1's type, i.e., initial fallback payoff, on the other hand, is 0, which is always publicly observable.

In terms of the structure of our game, it is a general ultimatum game with a Proposer (Player 1) and a Responder (Player 2), but there are two key features in our model that differentiates it from a standard ultimatum game. The first one is that the final fallback position (or final fallback payoff) of Player 2 is determined by her investment $a_j \geq 0$ for each type $j \in \{L, H\}$. This investment is costly. The cost is given by a strictly convex and strictly increasing function $C(\cdot)$ that is the same for both types and is publicly known under each information structure. The assumptions on the cost function are as follows:

Assumption 1. *$C(0) = C'(0) = 0 < C'(a)$, $C''(a)$ $\forall$ $a > 0$.*

We assume that depending on her final fallback position Player 2 may command offers from the market. If there are multiple offers from the market, Player 2 would clearly choose the maximum offer among them, in the absence of any offer from Player 1. We assume that the level of her final fallback position will correspond to the maximum offer Player 2 will receive from the market. E.g., if the low type of Player 1 obtains no investment, then the maximum offer she will receive is zero.

The second key feature that differentiates our model from a standard ultimatum game is that Player 1 can make an offer to Player 2 at any point, be it a preemptive offer or a counter offer. In the noisy-information case, he can also announce in the beginning of the game a menu of signal-contingent or other conditional preemptive offers or simply a counter offer after Player 2 makes her investment.

Overall, first, Player 2 observes her type $j \in \{L, H\}$ and thus the corresponding initial fallback payoff $d_j \in \{0, d_H\}$. Player 1 may or not observe Player 2's type and her investment, i.e., whether it is the *complete-information case* or not. If it is not a complete-information case (i.e., if he cannot observe her type or investment), Player 1 at least has a prior belief $g > 0$ that Player 2 is of the high type. If Player 1 does not obtain any further signals (to enable her to update her belief), then it is the *no- or hidden-information case*. If Player 1 cannot observe Player 2's type but obtains informative signals about Player 2's type and investment (on top of her prior belief), then it is the *incomplete- or noisy-information case.*

Let $\beta(\cdot)$ denote an offer that Player 1 can extend at any point after he finds out whether it is a complete- or no- or incomplete-information case.

The initial (zeroth) stage in timeline for our game for any information structure is such that Player 2 observes her type $j \in \{L, H\}$ and thus the corresponding initial fallback payoff $d_j \in \{0, d_H\}$, and both players find out the prevailing information structure. The remaining stages in the timeline depend on the information structure.

We use Perfect Bayesian Equilibrium (PBE) as the solution concept. PBE imposes sequential and Bayesian rationality. Bayesian rationality means that Player 1 uses Bayesian update rule whenever he observes an information about the initial or final fallback payoff of Player 2. Sequential rationality requires that the prescribed equilibrium behavior constitutes a Nash equilibrium in all of the subgames.

When Player 1 can make at most one offer, we consider the above three information structures in the next section. In the section following the next section, we will consider the multiple-offers case.

4 At Most One Offer

4.1 The Complete-information Case

The timeline for the complete information case that is beyond the initial, zeroth, stage, where players publicly discover the information structure, is as follows:

1. Player 1 finds out about Player 2's type, d_j.

2. Player 1 can make a preemptive offer $\beta(d_j)$ or wait for Player 2's investment and her market offers to make a counter offer to her.

   If a preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If the preemptive offer is rejected, then disagreement ensues where Player 1 receives 0, and Player 2 of type j chooses the investment level a_j to receive the maximum offer from the market with the payoff $d_j + a_j - C(a_j)$.

   If no preemptive offer is made, game proceeds to the next stage.

3. Player 2 chooses an investment level $a_j \geq 0$, which determines her maximum offer from the market.

4. Player 2's maximum offer is observed by both players. Player 1 can then choose to make a counter offer, $\beta(d_j + a_j)$. If a counter offer is made and accepted, it is enforced with its corresponding payoffs.

   If Player 1 does not make a counter offer or his counter offer is rejected, then disagreement ensues where Player 1 receives 0 and Player 2 of type j receives $d_j + a_j - C(a_j)$.

In this case PBE reduces to Subgame-Perfect Nash Equilibrium. Since each information set is going to be a singleton, for each possible offer one can simply calculate the equilibrium outcome through Backward Induction. Therefore, Player 1's preemptive offer at the second stage, $\beta(d_j)$, depends on what he expects about Player 2's investment in later stages.

If Player 2 of type j were to expect a counter offer, she would make an investment to maximize $d_j + a_j - C(a_j)$. The optimal investment in that case would be given by the following first-order condition

$$1 = C'(a^*) \Leftrightarrow a^* = (C')^{-1}(1), \tag{1}$$

If $d_H + a^* \geq 1$, then type H would invest by $1 - d$ and is trivially able to claim the whole bargaining surplus. Thus, we will focus on the more interesting case, in which $d_H + a^* < 1$.

Assumption 2. *$d_H + a^* < 1$, where a^* is defined in equation* (1).

With Assumption 2 above, we can expect that there is always agreement, since Player 1 would tailor his counter offer according to Player 2's maximum market offer and expect to get a payoff that is higher than his own fallback payoff, which is 0.

Now, given the expected counter-offer and Player 2's investment level, Player 2 of type j will get $a^* + d_j - C(a^*)$ if the game proceeds to the stage where Player 2 gets to make her investment. Then, if Player 1 can simply make a preemptive offer that promises $a^* + d_j - C(a^*)$ to Player 2, he can save $C(a^*)$ since Player 2 will not make any investment. This calculation establishes the Proposition 1 below.

Theorem 1. *Suppose that Assumption 2 holds. In the complete-information case with at most one offer, there is a unique equilibrium, in which Player* 1 *makes the preemptive offer $\beta(d_j) = a^* + d_j - C(a^*)$. Player* 2 *of each type accepts this preemptive offer and consequently does not make any investment, leading to an efficient outcome.*

4.2 The No-information Case

In this section, we consider the other extreme case, where player 1 cannot not receive any information about the type and investment of Player 2 before her maximum market offer is realized. Therefore, if player 1 considers making a preemptive offer he can only rely on his prior beliefs about the type of Player 2. The remainder of the timeline for the no-information case is as follows:

1. Player 1 finds out his prior belief, g, regarding Player 2's type.

2. Player 1 can make a preemptive offer $\beta(g)$ or wait for Player 2's investment and her market offers to make a counter offer to her.

   If a preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If the preemptive offer is rejected, then disagreement ensues where Player 1 receives 0, and Player 2 of type j chooses the investment level a_j to receive the maximum offer from the market with the payoff $d_j + a_j - C(a_j)$.

   If no preemptive offer is made, game proceeds to the next stage.

3. Player 2 chooses an investment level $a_j \geq 0$, which determines her maximum offer from the market.

4. Player 2's maximum offer is observed by both players. Player 1 can then choose to make a counter offer, $\beta(d_j + a_j)$.

   If a counter offer is made and accepted, it is enforced with its corresponding payoffs.

   If Player 1 does not make a counter offer or his counter offer is rejected, then disagreement ensues where Player 1 receives 0 and Player 2 of type j receives $d_j + a_j - C(a_j)$.

Player 1's preference over preemptive and counter offers depends on his prior belief about the type of Player 2. If Player 1 believes that Player 2 is of the high type with a very high likelihood, he would prefer to tailor his offer specifically for Player 2 of the high type. At the other extreme, Player 1 would target only the low type if he believes to a high degree that Player 2 is of the low type. In the intermediate cases though, Player 1's preferences over preemptive and counter offers depend on how costly Player 1's investment is and thus on how much Player 1 can gain by making a high enough offer to incentivize Player 2 not to invest. Theorem 2 below provides the conditions on the prior belief and the investment costs that determines what kind of offer Player 1 will choose.

Theorem 2. *Suppose that Assumption 2 holds. In the no-information case with at most one offer, there is a unique equilibrium in each case below:*

1. If

$$\frac{C(a^*)}{1-a^*-d_H+C(a^*)} < \frac{d_H}{1-a^*+C(a^*)} < \frac{d_H-C(a^*)}{d_H}, \tag{2}$$

(a) *Player 1 makes the high preemptive offer* $a^*+d_H-C(a^*)$ *if* $g \geq \frac{d_H-C(a^*)}{d_H}$. *Both types of Player 2 accept this offer and not invest;*

(b) *Player 1 waits for Player 2's investment and then makes the counter offer if* $\frac{C(a^*)}{1-a^*-d_H+C(a^*)} < g < \frac{d_H-C(a^*)}{d_H}$. *Both types of Player 2 accept this offer after making their investment,* a^**;*

(c) *Player 1 makes the low preemptive offer* $a^*-C(a^*)$ *if* $g \leq \frac{C(a^*)}{1-a^*-d_H+C(a^*)}$. *Only the low type of Player 2 accepts this offer and not invest, while the high type of Player 2 rejects it and invests by* a^*.

2. If

$$\frac{d_H-C(a^*)}{d_H} < \frac{d_H}{1-a^*+C(a^*)} < \frac{C(a^*)}{1-a^*-d_H+C(a^*)}, \tag{3}$$

(a) *Player 1 never makes a counter offer to Player 2.*

(b) *Player 1 makes the high preemptive offer* $a^*+d_H-C(a^*)$ *if and only if* $g \geq \frac{C(a^*)}{1-a^*-d_H+C(a^*)}$. *Player 1 makes the low preemptive offer* $a^*-C(a^*)$ *otherwise.*

Proof. If player 1 waits for the investment of Player 2 and makes the matching counter offer, Player 2 would invest as much as she can to maximize her bargaining position. Therefore, Player 1 would expect to receive the following ex-ante payoff:

$$(1-a^*)(1-g)+(1-a^*-d_H)g.$$

If Player 1 chooses to make a preemptive offer, there are only two viable options the optimal disagreement payoff of the low type or the high type. Then, the expected payoff for the preemptive low offer is

$$(1-a^*+C(a^*))(1-g),$$

while the expected payoff for the preemptive high offer is

$$1 - a^* - d_H + C(a^*).$$

Now, counter-offer is better for player 1 than preemptive low offer if and only if

$$(1 - a^*)(1 - g) + (1 - a^* - d_H)g > (1 - a^* + C(a^*))(1 - g) \Leftrightarrow$$
$$g > \frac{C(a^*)}{1 - a^* - d_H + C(a^*)}.$$

Counter offer is better for player 1 than preemptive high offer if and only if

$$(1 - a^*)(1 - g) + (1 - a^* - d_H)g > 1 - a^* - d_H + C(a^*) \Leftrightarrow$$
$$g < \frac{d_H - C(a^*)}{d_H}.$$

Note that if $d_H \leq C(a^*)$, preemptive high offer is always better than counter offer. Finally, preemptive high offer is better than preemptive low offer if and only if

$$(1 - a^* + C(a^*))(1 - g) \leq 1 - a^* - d_H + C(a^*) \Leftrightarrow$$
$$g \geq \frac{d_H}{1 - a^* + C(a^*)},$$

which is less than 1 since $d_H + a^* < 1$.

The comparisons of the types of offers gave us three thresholds for the prior probability g in terms of the cost function and the initial bargaining position of Player 2 of high type. Relative positions of these thresholds give us a map of Player 1's choices over the ranges of the prior probability. Since there are three thresholds, there are six orderings of these thresholds to consider. Four of these cases lead to contradiction. To see this consider the following two couples of comparisons:

1. If

$$\frac{d_H}{1 - a^* + C(a^*)} < \min\left\{\frac{d_H - C(a^*)}{d_H}, \frac{C(a^*)}{1 - a^* - d_H + C(a^*)}\right\},$$

we arrive at a contradiction, since when

$$\frac{d_H}{1-a^*+C(a^*)} < g < \min\left\{\frac{d_H - C(a^*)}{d_H}, \frac{C(a^*)}{1-a^*-d_H+C(a^*)}\right\},$$

preemptive low offer is better than counter offer, while counter offer is better than preemptive high offer, which is better than preemptive low offer.

2. If

$$\max\left\{\frac{d_H - C(a^*)}{d_H}, \frac{C(a^*)}{1-a^*-d_H+C(a^*)}\right\} < \frac{d_H}{1-a^*+C(a^*)},$$

we arrive at a contradiction, since when

$$\max\left\{\frac{d_H - C(a^*)}{d_H}, \frac{C(a^*)}{1-a^*-d_H+C(a^*)}\right\} < g < \frac{d_H}{1-a^*+C(a^*)},$$

counter offer is better than preemptive low offer, which is higher than preemptive high offer, which is in turn higher than counter offer.

The remaining two cases are described as in the hypothesis.

□

Switching from preemptive high offer $a^* + d_H - C(a^*)$ to the counter offer provides an expected net gain of $d_H(1-g) - C(a^*)$. Intuitively, the counter offer enables Player 1 to avoid allocating too much surplus to the low type, which is represented by the positive term $d_H(1-g)$ in the net gain, but also there is the inefficiency due to investment, which is represented by the negative term $-C(a^*)$. Therefore, optimality of a counter offers relies on the relation between the initial fallback position d_H and the cost of the disagreement investment $C(a^*)$. Note that it does not necessarily imply that $C(a^*)$ increases if the cost function shifts up. Therefore, the comparative statics on the optimality of the counter offers may not be monotonic.

To illustrate the conditions (2) and (3) that determine whether a counter offer could ever be preferred by Player 1, we parametrize the model by assuming a quadratic cost function such that $C(a) = ka^2$ for any investment level $a \geq 0$. Then, condition (2) can be expressed as

$$1 \geq 4k(1+d_H) - 16k^2 d_H(1-d_H), \tag{4}$$

after substituting the disagreement investment $a^* = \frac{1}{2k}$ and the cost of investment $C(a^*) = \frac{1}{4k}$. Note that, as the cost parameter k increases, the cost of investment falls

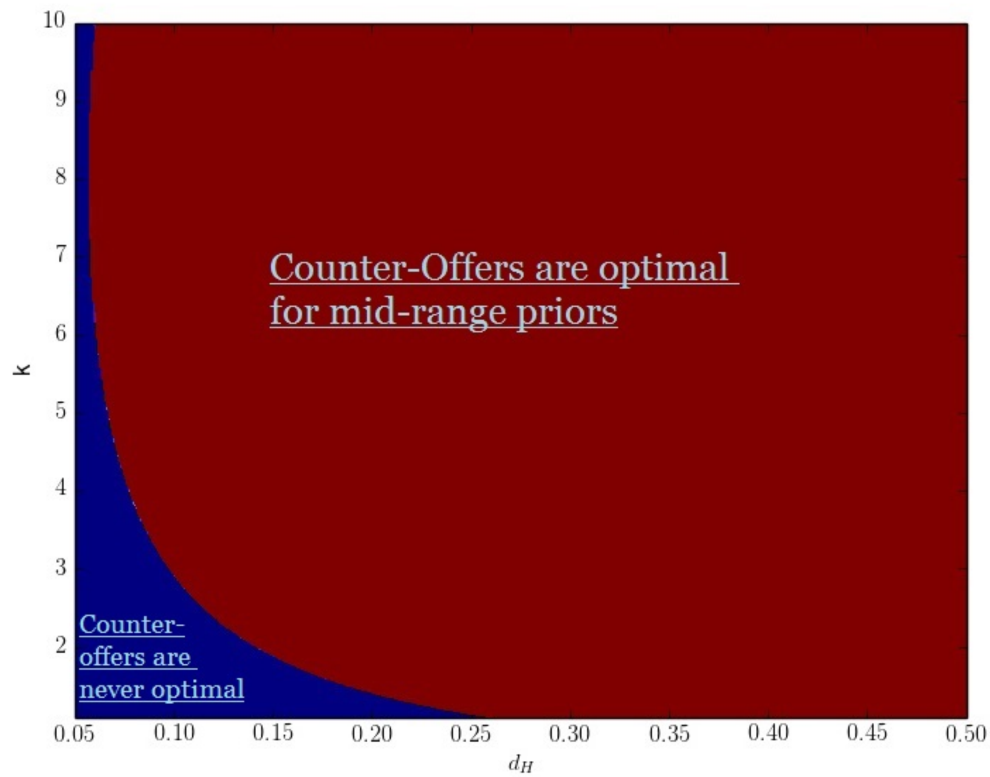

as well, since the investment level a^* falls more rapidly than the rate that the cost increases. This implies that as the cost of investment increases, it becomes more likely that the net gain of the counter offer $d_H(1-g) - C(a^*)$ increases further. Figure 4.2 confirms our intuition that for quadratic cost functions, higher cost of investment makes it easier for a counter offer to be optimal. The red-shaded area shows the range of parameter couples (d_H, k) that make it possible for counter offer to be ever preferred by Player 1.

4.3 The Noisy-information Case

In this section, we consider an intermediate case compared to the two extreme cases we have considered above. After the investment decision of Player 2, Player 1 receives two separate signals about Player 2. Signal $\tau = d_j + \sigma_d \varepsilon_d$ is a noisy signal about the type of Player 2, where σ_d is a signal dispersion parameter and ε_d is the random noise. As $\sigma_d$approaches to 0, the signal τ becomes a more precise signal about the type of Player 2. The random noise ε_d is distributed with respect to a probability distribution, of which the probability distribution function is $f_d(\cdot)$. We assume that $f_d(\cdot)$ is of full support on the Real numbers, continuously differentiable, and log-concave; that is, $\frac{f'(\tau)}{f(\tau)}$ is strictly decreasing in signal τ. Since log-concavity implies monotone-likelihood ratio property, the higher the signal τ the higher the likelihood that Player 2 is of the high type. The second signal α is about the investment of Player 2. It is as follows:

$$\alpha = \begin{cases} a_j + \sigma_a \varepsilon_a & \text{if } a_j > 0 \\ \emptyset & \text{if } a_j = 0, \end{cases}$$

where a_j is the investment by Player 2 of type j. We assume that if Player 2 does not make any investment, it is impossible for Player 1 to observe any positive signal. This captures the idea that an administrator may observe some activities carried out by the employee if employee is actively searching for a better offer in the market, but it is impossible for the administrator to observe anything if the employee does not engage in any investment regarding her fallback position. Here too σ_a is the dispersion parameter, this time corresponding to the investment signal. The random noise ε_a is distributed with respect to a probability distribution, of which the probability distribution function is $f_a(\cdot)$. We again assume that $f_a(\cdot)$ is of full support on the Real numbers, continuously differentiable, and log-concave. We assume that the two random noise variables are not correlated.

The remainder of the timeline for the noisy-information case is as follows:

1. Player 1 finds out his prior belief, g, regarding Player 2's type.

2. Player 1 can make a preemptive offer $\beta(g)$.

   If a preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If the preemptive offer is rejected, then disagreement ensues where Player 1 receives 0, and Player 2 of type j chooses the investment level a_j to receive the maximum offer from the market with the payoff $d_j + a_j - C(a_j)$.

   If no preemptive offer is made, game proceeds to the next stage.

3. Player 1 receives a noisy signal, τ, about Player 2's type.

4. Player 1 can make a signal-contingent preemptive offer $\beta(g, \tau)$.

   If a signal-contingent preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

If a signal-contingent preemptive offer is rejected, then disagreement ensues where Player 1 receives 0, and Player 2 of type j chooses the investment level a_j to receive the maximum offer from the market with the payoff $d_j + a_j - C(a_j)$.

If no signal-contingent preemptive offer is made, the game proceeds to the next stage.

5. Player 2 chooses an investment level $a_j \geq 0$, which determines her maximum offer from the market.

6. Player 1 receives a signal, α, about Player 2's investment level such that either $\alpha = \emptyset$ or $\alpha \in \mathbb{R}_+$.

7. Then Player 1 can make a double signal-contingent offer $\beta(g, \tau, \alpha)$.

   If a double signal-contingent offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If the double signal-contingent preemptive offer is rejected, then disagreement ensues where Player 1 receives 0, and Player 2 of type j chooses the investment level a_j to receive the maximum offer from the market with the payoff $d_j + a_j - C(a_j)$.

   If no double signal-contingent preemptive offer is made, the game proceeds to the next stage.

8. Player 2's maximum offer is observed by both players. Player 1 can then choose to make a counter offer, $\beta(d_j + a_j)$.

   If a counter offer is made and accepted, it is enforced with its corresponding payoffs.

   If Player 1 does not make a counter offer or his counter offer is rejected, then disagreement ensues where Player 1 receives 0 and Player 2 of type j receives $d_j + a_j - C(a_j)$.

As before, Player 1 chooses to make an offer that maximizes the chance of agreement without giving too much surplus to Player 2. It would be optimal for Player 1 to offer the disagreement surplus $a^* + d_j - C(a^*)$ so that Player 2 of type j does not make any investment and accepts the corresponding offer. However, since the information that Player 1 expects to receive is not fully informative, Player 1 cannot always guarantee an agreement unless he either makes the preemptive high offer $a^* + d_H - C(a^*)$ and thus gives too much surplus to the low type or he ignores the signals and allows Player 2 to make investment where he would match her final fallback position (and thus the maximum offer she receives from the market) with a counter offer.

We first consider the signal-contingent preemptive offer $\beta(g, \tau)$ that is conditional only on the signal τ about the type of Player 2. With that preemptive offer Player 1 seeks to prevent Player 2 from making any investment.

Proposition 1 describes the signal-contingent offer that kills any incentives of Player 2 to insure herself against disagreement or manipulate the signal that Player 1 will receive.

Proposition 1. *In the noisy-information case with at most one offer, Player 1's signal-contingent preemptive offer $\beta(g, \tau)$ is described as follows*

$$\beta(g,\tau) = \begin{cases} a^* - C(a^*) & \text{if } \tau \leq \bar{\tau} \\ a^* + d_H - C(a^*) & \text{if } \tau > \bar{\tau}, \end{cases}$$

where the signaling threshold is uniquely defined by the following indifference condition:

$$(1-g)d_H f_d\left(\frac{\bar{\tau}}{\sigma_d}\right) = g(1 - a^* - d_H + C(a^*))f_d\left(\frac{\bar{\tau} - d_H}{\sigma_d}\right) \tag{5}$$

Proof. Since Player 2 can be of two types, it is sufficient for Player 1 to include only two offers. To see this, note that if Player 1 chooses two offers that Player 2 of high type would accept, Player 1 can always pick the lower offer and still guarantee agreement. Similarly, if Player 1 chooses low offers that the low type would accept and the high type would reject, then again Player 1 would choose the lowest offer possible.

Given any two offers $\beta_L < \beta_H$, Player 1 compares the posterior expected payoff given the signal he receives. When Player 1's posterior belief that Player 2 is of the high type is low enough, Player 1 would make an offer that only the low type would accept. Then, as in the analysis in the previous sections, Player 1 would choose to offer $a^* - C(a^*)$ that will guarantee an agreement with the low type and keep her from making an investment. If the posterior belief of Player 1 assigns enough likelihood to the high type, Player 1 would target the high type and offer $a^* + d_H - C(a^*)$.

Log-concavity implies that the posterior probability that the type of Player 2 is high increases with the signal. Then, there exists a unique signal $\bar{\tau}$ that makes Player 1 indifferent between choosing the high or the low offer. This indifference condition is given as

$$\frac{(1-a^*+C(a^*))(1-g)f_d\left(\frac{\bar{\tau}}{\sigma_d}\right)}{(1-g)f_d\left(\frac{\bar{\tau}}{\sigma_d}\right)+gf_d\left(\frac{\bar{\tau}-d_H}{\sigma_d}\right)} = 1-a^*-d_H+C(a^*),$$

since the high offer is accepted by both types but the low offer is accepted only by the low type. Re-arranging the indifference condition above gives the equation (5).

□

Player 1 makes a signal-contingent preemptive offer only if he expects to receive a payoff higher than what a signal-noncontingent preemptive offer could generate for Player 1. The payoff that Player 1 expects to receive from a signal-contingent offer before he receives the signal is

$$\begin{aligned}(1-g)\left(F_d\left(\frac{\bar{\tau}}{\sigma_d}\right)(1-a^*+C(a^*))+\left(1-F_d\left(\frac{\bar{\tau}}{\sigma_d}\right)\right)(1-a^*-d_H+C(a^*))\right)\\ +g\left(\left(1-F_d\left(\frac{\bar{\tau}-d_H}{\sigma_d}\right)\right)(1-a^*-d_H+C(a^*))\right) \quad (6)\end{aligned}$$

This ex-ante payoff depends on how informative the signal is, i.e., on the value of the signal dispersion parameter σ_d. As the signal gets more precise, Player 1 can rely more on the signal to make the offer tailored for the correct type of Player 2. This, in principle, should increase the expected payoff for a signal-contingent offer. Proposition 2 below states that the signal-contingent offers have an advantage over signal-noncontingent preemptive offers, regardless of the precision level of the signal.

Proposition 2. *Suppose that Assumption 2 holds. In the noisy-information case with at most one offer, the optimal signal-contingent preemptive offer yields Player 1 a higher ex-ante payoff than both the signal-noncontingent preemptive high offer and the signal-noncontingent preemptive low offer.*

Proof. Player 1's expected payoff from the (optimal) signal-contingent preemptive offer (see Equation 6) is $(1-g)(1-a^*+C(a^*)-(1-F_{d0})d_H)+g(1-F_{dH})(1-a^*-d_H+C(a^*))$, where $F_0 = Pr(\tau \leq \bar{\tau}|j=L) = F_d\left(\frac{\bar{\tau}}{\sigma_d}\right)$, $F_{dH} = Pr(\tau \leq \bar{\tau}|j=H) = F_d\left(\frac{\bar{\tau}-d_H}{\sigma_d}\right)$.

Note that in the (optimal) signal-contingent preemptive offer, the signal threshold is given by equation 5. If Player 1 chooses a different signal threshold in the signal-contingent preemptive offer, then Player 1 must obtain a lower expected payoff. Note that (i) if the signal threshold that Player 1 chooses is $+\infty$, then the signal-contingent preemptive offer is always $a^* - C(a^*)$ (regardless of τ), and thus Player 1 will obtain the same expected payoff as the signal-noncontingent preemptive low offer, and (ii) if the signal threshold that Player 1 chooses is $-\infty$, then the signal-contingent preemptive offer is always $a^* + d_H - C(a^*)$ (regardless of τ), and thus Player 1 will obtain the same expected payoff as the signal-noncontingent preemptive high offer. So, the (optimal) signal-contingent preemptive offer must yield Player 1 a higher expected payoff than both the signal-noncontingent preemptive high offer and the signal-noncontingent preemptive low offer.

□

Proposition 2 provides a comparison between the two types of preemptive offers, signal-contingent ones and signal-noncontingent ones.

The next step in the analysis of the case of noisy signals is to compare the two types of "post-investment" offers: the double signal-contingent offers, $\beta(\tau, \alpha)$, and the counter offer $\beta(a_j + d_j)$ for each type $j \in \{L, H\}$. Proposition 3 below states that making a counter offer is the only equilibrium outcome in the subgame where Player 2 invests to insure herself against disagreement.

Proposition 3. *Suppose that Assumption 2 holds. In the noisy-information case, if Player* 1 *does not make any preemptive offer and thus Player* 2 *makes an investment, Player* 1 *chooses to wait for the realization of the maximum offer that Player* 2 *can generate and to match that offer with a counter offer. Player* 2*, expecting to receive a counter offer, invests by* a^*.

Proof. Suppose that Player 2 of type j makes an investment at level a_j such that $0 < a_L < a_H + d_H$. We will verify this supposition later on. Then, the expected payoff for double signal-contingent offer for any couple of signals (τ, α) is

$$\max\{P(j = L|\tau, \alpha)(1 - a_L), 1 - a_H - d_H\},$$

while the expected payoff for the counter-offer is

$$P(j = L|\tau, \alpha)(1 - a_L) + (1 - P(j = L|\tau, \alpha))(1 - a_H - d_H),$$

which is clearly greater than the expected payoff for the double signal-contingent offer since $a_L < a_H + d_H$.

Now, given that player 1 always prefers counter-offers, the optimal investment for Player 2 of either type is a^*. Note that such an investment behavior by Player verifies our supposition above.

□

In addition, the only viable options Player 1 has in equilibrium are signal-contingent and counter offers. Further, there is a particular cutoff between these two options.

Theorem 3. *Suppose that Assumption 2 holds. In any equilibrium in the noisy-information case with at most one offer, Player* 1 *will never make a signal-noncontingent preemptive offer or a double signal-contingent offer.*

Let τ *be the signal realization of Player* 2*'s type, and*

$$g(\tau) = \frac{g f_d\left(\frac{\tau - d_H}{\sigma_d}\right)}{(1-g) f_d\left(\frac{\tau}{\sigma_d}\right) + g f_d\left(\frac{\tau - d_H}{\sigma_d}\right)} \tag{7}$$

be Player 1*'s posterior about Player* 2*'s type after Player* 1 *receives the signal realization* τ*. Then, Player* 1 *chooses a counter offer over signal contingent offer if and only if*

$$\frac{C(a^*)}{1 - a^* - d_H + C(a^*)} < g(\tau) < \frac{d_H - C(a^*)}{d_H}. \tag{8}$$

Proof. By Proposition 2, we know that whenever a signal-noncontingent offer is better for Player 1 than counter-offer, signal-contingent offer $\beta(\tau)$ should also better than counter-offer. Then, by Theorem 2 the only case that counter-offer is better than signal-noncontingent offers is the one that is laid out in the hypothesis.

To see that counter-offer is better for Player 1 than the signal-contingent offer in this case, note that signal-contingent offer promises the same payoff as the low signal-noncontingent offer if $\tau < \bar{\tau}$, and the same payoff as the high signal-noncontingent offer if $\tau \geq \bar{\tau}$. Since, counter-offer promises a higher payoff than both of the signal-noncontingent offers in this case, it also promises a higher payoff than the signal-contingent offer $\beta(\tau)$.

□

Recall from our previous discussion following our main result in the no-information case that, for quadratic cost functions, higher cost of investment makes it easier for a counter offer to be optimal. The same intuition holds for the noisy-information case as well. When the posterior belief of Player 1 is not decisive enough, he chooses to wait for the counter-offer opportunity. The main difference between the no- and noisy-information cases on the optimality of counter offer is that Player 1's posterior belief after observing the signal τ is more precise than his prior belief. Therefore, in the noisy information case it is harder for Player 1 to remain too uncertain about the type of Player 2 to wait for the counter-offer opportunity.

5 Multiple Offers

In the previous section, we have assumed so far that whenever Player 2 rejects an offer of Player 1, the game proceeds to the disagreement stage. However, it might also be possible in some contexts that Player 1 is able to amend his offer to renegotiate before Player 2 receives her final fallback payoff. In addition, theoretically it is always a relevant attempt to consider a setup without any restrictions on the number of offers that can be made. Consequently in this section we consider an alternative timeline where game does not proceed to disagreement stage when Player 2 rejects a preemptive offer. The main intuitive implication of this alternative timeline is that Player 1 can strategically choose to start by making a low preemptive offer first to allow for gradual increases in his offers at the later stages of the game (this should not be explicitly costly for him since he does not face any discounting).

Since multiple offers possibility would not change the main result of the complete-information case, here we will confine ourselves to the cases of no-information and noisy information.

For simplicity, we will assume that preparation of an offer by Player 1 and considering to accept or reject an offer by Player 2 may take time, and therefore Player 1 can make a maximum finite number of offers m at any segment of the game.

5.1 The No-information Case

The timeline that we will consider in this section for the no-information case is as follows:

1. Player 1 finds out his prior belief, g, regarding Player 2's type.

2. Player 1 can make any finite number, k, of consecutive preemptive offers $(\beta_k(g))_{k=1,m}$ or wait for Player 2's investment and her market offers to make a counter offer to her.

   If a preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If no preemptive offer is made or if all of the preemptive offers are rejected, game proceeds to the next stage.

3. Player 2 chooses an investment level $a_j \geq 0$, which determines her maximum offer from the market.

4. Player 2's maximum offer is observed by both players. Player 1 can then choose to make any finite number, k, of consecutive counter offers $(\beta_k(d_j + a_j))_{k=1,m}$.

   If a counter offer is made and accepted, it is enforced with its corresponding payoffs.

   If Player 1 does not make any counter offers or all of his counter offers are rejected, then disagreement ensues where Player 1 receives 0 and Player 2 of type j receives $d_j + a_j - C(a_j)$.

Theorem 4 below shows that Player 1 prefers to combine the low preemptive offer $a^* - C(a^*)$ with the high counter offer $a^* + d_H$ for a relatively large range of prior beliefs. The downside of this strategy for Player 1 is that he cannot capture the cost of investment if Player 2 is of the high type. Therefore, when the cost of investment is exceedingly high and Player 1 is confident enough that Player 2 is of high type, he chooses to make the high preemptive offer to guarantee an agreement without any investment.

Theorem 4. *In the no-information case with multiple offers, Player* 1 *uses one of the two following strategies depending on his priors. (1) If* $g \leq \frac{d_H}{d_H + C(a^*)}$*, Player* 1 *makes the low preemptive offer* $a^* - C(a^*)$*, and if rejected by the high type of Player* 2*, he makes the high counter-offer* $a^* + d_H$ *after the investment of Player* 2 *which is accepted. (2) If* $g > \frac{d_H}{d_H + C(a^*)}$*, Player* 1 *makes the high preemptive offer, which results in agreement with any type of Player* 2 *without any investment.*

Proof. Since making an offer at the first stage is costless, Player 1 would always make an offer that at least Player 2 of low type would accept. Recall that minimum such offer is $a^* - C(a^*)$. Suppose for now that Player 1 is allowed to do only one preemptive offer.

If Player 1 makes a low offer, the expected payoff is $(1-g)(1-a^*+C(a^*)) + g(1-a^*-d_H)$, since if Player 2 is of high type, she will reject the preemptive offer and proceed to investment, which results in the corresponding counter offer. If Player 1 makes a high preemptive offer $a^* + d_H - C(a^*)$, both types of Player 2 would immediately accept the offer. Then, the strategy of low preemptive and high counter offer is better than the high preemptive offer if and only if

$$(1-g)(1-a^*+C(a^*)) + g(1-a^*-d_H) \geq 1-a^*-d_H+C(a^*) \Leftrightarrow g \leq \frac{d_H}{d_H+C(a^*)}.$$

Now, we will show that even if Player 1 has access to a finite number of consecutive preemptive offers, he cannot do better than a single offer. We can divide the set of finite equilibrium preemptive offer strategies into the repeated relatively low preemptive offers that only the low type would accept and any finite sequence of offers that includes at least one high preemptive offer that both types would accept. It is clear that the repeated low offer promises the same payoff as the single low offer. All other strategies promise the same payoff as the single high offer. This because if Player 2 expects to receive a high preemptive offer in one of the future preemptive offers, she does not have any incentive to accept a low offer. Then, rejection to a low offer does not give any information to Player 1, which implies that once a finite offer strategy with at least one high offer is chosen, Player 1 does not have any incentive to change to another finite offer strategy after observing a rejection of a low offer.

□

Theorem 4 shows that since Player 1 combines the low preemptive offer with counter offer, the analysis of his equilibrium play reduces down to the comparison of two strategies, which leads to the single threshold, $\frac{d_H}{d_H+C(a^*)}$, for the prior beliefs.

5.2 The Noisy-information Case

When Player 1 has access to noisy signals, the possibilities for Player 1's strategy of offers substantially increase, since this gives rise to more possibilities of negotiation for Player 1. The timeline with multiple offers in the noisy information case is as follows:

1. Player 1 finds out his prior belief, g, regarding Player 2's type.

2. Player 1 can make any finite number, k, of consecutive preemptive offers $(\beta_k(g))_{k=1,m}$ or wait for Player 2's investment and her market offers to make a counter offer to her.

   If a preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If no preemptive offer is made or if all of the preemptive offers are rejected, game proceeds to the next stage.

3. Player 1 receives a noisy signal, τ, about Player 2's type.

4. Player 1 can make any finite number, k, of consecutive signal-contingent preemptive offers $(\beta_k(g, \tau))_{k=1,m}$.

   If a signal-contingent preemptive offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If no signal-contingent preemptive offer is made or if all of the signal-contingent preemptive offers are rejected, the game proceeds to the next stage.

5. Player 2 chooses an investment level $a_j \geq 0$, which determines her maximum offer from the market.

6. Player 1 receives a signal, α, about Player 2's investment level such that either $\alpha = \emptyset$ or $\alpha \in \mathbb{R}_+$.

7. Then Player 1 can make any number, k, of consecutive double signal-contingent offers $(\beta_k(g, \tau, \alpha))_{k=1,m}$.

   If a double signal-contingent offer is made and accepted, the game ends with the enforcement of the corresponding payoffs.

   If no double signal-contingent preemptive offer is made or if all of the double signal-contingent preemptive offers are rejected, the game proceeds to the next stage.

8. Player 2's maximum offer is observed by both players. Player 1 can then choose to make any number, k, of consecutive counter offers, $(\beta_k(d_j + a_j))_{k=1,m}$.

If a counter offer is made and accepted, it is enforced with its corresponding payoffs.

If Player 1 does not make any counter offers or all of his counter offers are rejected, then disagreement ensues where Player 1 receives 0 and Player 2 of type j receives $d_j + a_j - C(a_j)$.

As was the case in the preceding subsection as well, Theorem 5 below shows that Player 1 combining the low preemptive offer with the high counter offer is the optimal decision for him in the noisy information case as well. In this case, Player 1 makes his decisions contingent on the type signal instead of his priors: Player 1 makes the low preemptive offer if he observes a relatively low signal τ which only the low type of Player 2 accepts and the high type rejects. Upon rejection, Player 1 makes a the appropriate counter offer after observing the maximum offer Player 2 receives from the market. Player 1 makes the high preemptive offer if he observes a relatively high signal τ which both types of Player 2 accept without any investment.

Theorem 5. *In the case of noisy signals τ with multiple offers, Player 1 makes the low preemptive offer $a^* - C(a^*)$ after observing the signal τ if $\tau \leq \bar{\tau}$, where $\bar{\tau}$ is uniquely defined by the following equation*

$$g(\bar{\tau}) = \frac{g f_d\left(\frac{\bar{\tau}-d_H}{\sigma_d}\right)}{(1-g) f_d\left(\frac{\bar{\tau}}{\sigma_d}\right) + g f_D\left(\frac{\bar{\tau}-d_H}{\sigma_d}\right)} = \frac{d_H}{d_H + C(a^*)}.$$

Player 1 makes the high preemptive offer $a^* + d_H - C(a^*)$ after observing the signal if $\tau > \bar{\tau}$. Any type of Player 2 accepts the high preemptive offer without any investment, while only the low type of Player 2 accepts the preemptive low offer, in which case Player 2 of high type makes investment and accepts the corresponding counter offer.

Proof. First, note that the availability of multiple offers at any stage does not improve Player 1's payoff compared to single offer at each stage. Therefore, we assume that Player 1 can make at most one offer at every stage for the rest of the proof.

Similar to Proposition 3, Player 1 would always prefer to make a counter-offer instead of double-signal contingent post-investment offer. In that case, Player 1 would make the low counter-offer a^* before the final disagreement payoff is realized, since such an offer is costless for Player 1. Then, if he observes rejection, he makes the high counter-offer. This implies that once Player 2 makes an investment the expected payoff of Player 1 is same as in the benchmark model with single binding offer.

Next, we consider the pre-investment stages and claim that Player 1 cannot monitor the type of Player 2 by any strategy of offers. Such a monitoring requires the existence of two offers $\beta_0 \in [a^* - C(a^*), \beta(\tau, reject))$, where $\beta(\tau, reject) \geq a^* + d_H - C(a^*)$. When Player 1 makes the offer β_0, the low type of Player 2 should accept while the high type should reject it. Then, a rejection would be perfectly informative about the type of Player 2, leading him to make the higher offer $\beta(\tau, reject)$. However, since $\beta(\tau, reject)$ $> \beta_0$, low type could mimic the high type by rejecting β_0.

Then, Player 1 always makes the low preemptive offer unless his posterior beliefs after observing τ are convincing enough for Player 1 to make a high offer. Then, by Proposition 2 and Theorem 4, Player 1 makes a binding high preemptive offer if his posterior belief is higher than $\frac{d_H}{d_H + C(a^*)}$. That is, Player 1 makes the high preemptive offer if and only if he does not observe an acceptance before he observes the noisy signal τ, and $\tau > \bar{\tau}$, where $\bar{\tau}$ is uniquely defined as in the hypothesis.

□

6 Concluding Remarks

When Player 1 is allowed to make at most one offer, we showed that in the complete-information case, the preemptive offer is the optimal one for Player 1, which leads to an efficient outcome avoiding any investment by Player 2. In the no-information case, depending on the circumstances a preemptive offer or a counter offer (which involves an investment by Player 2) can be optimal. We note that a higher cost of investment makes it easier for a counter offer to be optimal when Player 1's beliefs about the type of Player 2 are not very decisive. The intuition about the cost is that, as the cost parameter k increases, the cost of investment falls as well, since the investment level a^* falls more rapidly than the rate that the cost increases. This implies that as the cost of investment increases, it becomes more likely that the net gain of the counter offer $d_H(1-g) - C(a^*)$ increases further. Otherwise, Player 1 chooses to make either a high preemptive offer (if Player 1 has relatively high beliefs about Player 2 being the high type) which both types accept or a low preemptive offer (if Player 1 has relatively low beliefs about Player 2 being the high type) which only the low type accepts, but never to make a counter offer.

In the noisy-information case when Player 1 can make at most one offer, there are two more possibilities that Player 1 can choose from: signal-contingent preemptive offers and double-signal-contingent preemptive offers, as two separate signals - one signal about the type of Player 2 and, if any, another one about the investment level of Player 2 - could be observed by Player 1. We show that double-signal-contingent preemptive offers are never optimal for Player 1. We also show that signal-contingent preemptive offers dominate signal-noncontingent preemptive offers. It turns out that Player 1 chooses a counter offer over the signal-contingent preemptive offer as the cost of Player 2's investment goes down.

When Player 1 is allowed to make multiple offers, then the analysis of the complete-information case does not change but in the other cases he can cherry-pick by making multiple offers to separate the two types of Player 2. In the no-information case, Player 1 prefers to combine the low preemptive offer with the high counter offer. For a relatively large range of relatively low prior beliefs that he is facing the high type of Player 2, Player 1 makes the low preemptive offer which only the low type of Player 2 accepts. If this offer is rejected, then he makes the appropriate high counter offer to Player 2 after she receives offer from the market. If, however, Player 1 has a sufficiently high prior belief and the cost of investment is high, to guarantee an agreement without any investment, Player 1 makes the high preemptive offer which both types of Player 2 accept. It turns out that Player 1 combining the low preemptive offer with the high counter offer is the optimal decision for him in the noisy-information case as well. Note, however, that regardless of whether Player 1 can make at most one offer or multiple offers, he will make his decisions contingent on the type signal he receives instead of making them contingent on his priors, which means that Player 1's posterior belief after observing the signal τ will neessarily be more precise than her prior belief. Consequently, in the noisy information case Player 1 will have a lower tendency to make a counter offer to Player 2.

References

[1] Abreu, D., and Gul, F. Bargaining and Reputation. *Econometrica 68*, 1 (2000), 85–117.

[2] Anbarci, N., Skaperdas, S., and Syropoulos, C. Comparing Bargaining Solutions in the Shadow of Conflict: How Norms against Threats Can Have Real Effects. *Journal of Economic Theory 106*, 1 (2002), 1–16.

[3] Atakan, A. E., and Ekmekci, M. Bargaining and Reputation in Search Markets. *The Review of Economic Studies 81*, 1 (2014), 1–29.

[4] Baliga, S., and Sjöström, T. Bargaining and War: A Review of Some Formal Models. *Korean Economic Review 29*, 2 (2013), 235.

[5] Fisher, R., and Ury, W. L. *Getting to Yes: Negotiating Agreement Without Giving in.* New York: Penguin Books, 1981.

[6] Gennaioli, N., and Voth, H.-J. State Capacity and Military Conflict. *Review of Economic Studies 82*, 4 (2015), 1409–1448.

[7] Grossman, S. J., and Hart, O. D. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of political economy 94*, 4 (1986), 691–719.

[8] Hart, O., and Moore, J. Property Rights and the Nature of the Firm. *Journal of political economy 98*, 6 (1990), 1119–1158.

[9] Malhotra, D., and Gino, F. The Pursuit of Power Corrupts: How Investing in Outside Options Motivates Opportunism in Relationships. *Administrative Science Quarterly 56*, 4 (2011), 559–592.

[10] Morita, H., and Servátka, M. Group Identity and Relation-Specific Investment: An Experimental Investigation. *European Economic Review 58* (2013), 95–109.

[11] Myerson, R. B. *Game Theory: Analysis of Conflict.* Cambridge: Harvard University Press, 1991.

[12] Pinkley, R. L. Impact of Knowledge Regarding Alternatives to Settlement in Dyadic Negotiations: Whose Knowledge Counts? *Journal of Applied Psychology 80*, 3 (1995), 403–417.

[13] Pinkley, R. L., Neale, M. A., and Bennett, R. J. The Impact of Alternatives to Settlement in Dyadic Negotiation. *Organizational Behavior and Human Decision Processes 57*, 1 (1994), 97–116.

[14] Ramsay, K. W. Information, Uncertainty, and War. *Annual Review of Political Science 20* (2017), 505–527.

[15] Williamson, O. E. The Theory of the Firm as Governance Structure: From Choice to Contract. *The Journal of Economic Perspectives 16*, 3 (2002), 171–195.