USC FBE FINANCE SEMINAR
presented by John Zhu
FRIDAY, Sept. 23, 2016
10:30 am – 12:00 pm, **Room: JFF-414**

Renegotiation of Dynamically Incomplete Contracts

John Y. Zhu[1]

September 16, 2016

Abstract

I provide a micro-foundation for dynamically incomplete contracts that are renegotiated over time. The micro-foundation is based on showing that such contracts implement the optimal complete contract in a general dynamic financial contracting model provided the players have "preference-for-robustness." Preference-for-robustness is a class of dynamic max-min preferences defined in a setting where players have a fuzzy idea about events in the future. The paper culminates in an analysis of contracting under asymmetric information: The optimal dynamic contract under preference-for-robustness is shown to be a debt contract featuring state-contingent allocation of control rights and a refinance option.

JEL Codes: D86, G32, G33
Keywords: renegotiation, incomplete contracts, security design, robustness, financial contracting, optimal contracting, debt, control rights, refinance

[1]zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Sebastian Pfeil (discussant), Tomasz Piskorski (discussant), Alejandro Rivera (discussant), Xi Weng (discussant), Zenan Wu (discussant), Richard Kihlstrom, Daniel Ferreira, Alex Frankel, Andrew Postlewaite, and Michael Roberts for their advice and comments. I thank the seminar participants at Wharton, University of Michigan Ross, Swiss Finance Institute/HEC Lausanne, 2015 Workshop on Dynamic Incentive Problems - University of Bonn, 2016 Midwest Economic Theory Conference - University of Rochester, University of Pennsylvania, Finance Theory Group 2016 Spring Meetings - Boston College, 2016 University of Minnesota Corporate Finance Conference, 2016 SFS Cavalcade - University of Toronto, Tsinghua SEM, Peking University Applied Economics, Cornell, NYU.

1 Introduction

Consider the following scenario. An entrepreneur has a project and seeks financing. He secures a one year contract from a financier. Despite the one year window, the project is long-lived and will likely be productive and require further financing a year from now. Both parties understand this and know that a year from now, the maturing contract could be renegotiated, perhaps rolled over to another one year contract. Despite this understanding, the current contract does not explicitly lay out a state-contingent renegotiation plan for next year (e.g. some detailed set of instructions that is designed to ensure each party receives half of the surplus generated from renegotiation). Instead, if it turns out that next year there is some surplus to be captured by extending the relationship, then they will hammer out the details at that time. Furthermore, between now and next year, things might change, so that the parties will likely end up rewriting the contract before it even expires next year. Again, the parties do not pre-specify such an interim renegotiation, agreeing, instead, to cross that bridge if and when they get there.[2]

I call such a contract together with the example renegotiation paths a dynamically incomplete contract being renegotiated over time. My aim is to provide micro-foundations for these contracts. I do this by considering optimal complete contracting in a general dynamic financial contracting model and showing that every Pareto-optimal payoff can also be attained by a dynamically incomplete contract being renegotiated over time *if the players have "preference-for-robustness."* Preference-for-robustness is a class of weakly time-consistent preferences I introduce in the paper that models a decision maker who has a perpetually fuzzy idea as time passes about events in the future and is averse to this fuzziness. A decision maker with preference-for-robustness behaves in a way not unlike an ambiguity-averse decision maker. Indeed, preference-for-robustness is a simple way to bring that type of max-min decision making from the one-shot to the repeated setting, allowing it to be fruitfully applied to optimal dynamic contracting.

As the basic implementation results are gradually established, I build up to the key application: I introduce an asymmetric information problem into the model by assuming that the entrepreneur privately observes the contract relevant state of the world. I then show the optimal contract can be implemented by a dynamically incomplete debt contract that can be renegotiated and features a state-contingent allocation of control rights. This debt contract differs in key ways from many debt contracts derived in previous security design models featuring complete contracts. Typically, in such a model, when a contract is described as debt, what is meant is that the contract is static with a wedge payoff function of the form $D \wedge v$ where D is a constant and v is the random value of the underlying asset. In my model with asymmetric information, the optimal complete contract that looks like debt is more than this. The implementation of the complete contract specifies an initial short-term incomplete contract with the familiar wedge payoff. However, this payoff is preliminary and subject to renegotiation. Essentially, the wedge payoff sets the stage for renegotiation

[2]Empirical studies of privately placed debt show that they are renegotiated frequently and that typically these renegotiations are not planned out in advance. For example, Roberts (2015) finds that the average loan has a maturity of four and a half years and is renegotiated five times, but less than 28% of the renegotiations are triggered by an actual or anticipated violation of a pre-specified covenant.

by serving as the outside option in the ex-post bargaining game. Just as importantly, the implementation allocates bargaining power based on the state. If the project is in technical default - that is, if the current capital level is below D - then the financier gets to make a take-it-or-leave-it offer. Otherwise, the entrepreneur makes the take-it-or-leave-it offer.

The additional control rights facet of the debt contract is aligned with the incomplete contracting literature's perspective on debt. See, for example, Aghion and Bolton (1992). My work can be seen as providing support for this perspective by showing how a theory of optimal control rights allocation can emerge without assuming a particular incompleteness of the contracting space.

I emphasize that the choice of a state-contingent allocation of control rights does not have any intrinsic payoff impact ex-ante (for players who have preference-for-robustness): If there was no asymmetric information or if the entrepreneur could be forced to always tell the truth, then the debt contract's particular allocation of control rights can be changed to any other, and the ex-ante value of the contract would be unaffected. Allocating control rights specifically in the way that is done in a debt contract matters purely for incentive-compatibility. For example, I show that if, instead, the entrepreneur gets to make a take-it-or-leave-it offer when in technical default then the resulting misreporting incentives are strong enough to make the "debt" contract worthless to the financier.

The paper is related to various strands of the contracting literature. Were it not for the preference-for-robustness component, the optimal contracting problem would be a standard expected utility optimal dynamic contracting problem over complete contracts. By moving from expected utility to preference-for-robustness I show how the typically complex optimal dynamic contract reduces to a simple one that captures the essence of many real-life dynamic financing arrangements.[3]

My paper is also related to the effort to micro-found incomplete contracts, such as the null contract, in the hold-up model. See Hart and Moore (1999) and Segal (1999). The setting is one with observable but unverifiable signals. Since signals are not contractible, contracts can only be written over contractible messages sent by players. The papers show that when the environment is sufficiently "complex," the optimal renegotiation-proof contract over messages is no better than the null contract. Mukerji (1998), also looking at the hold-up model, points out that the observable but unverifiable assumption is not even needed provided players are sufficiently ambiguity averse. By looking at complete contracts under max-min decision making, Mukerji (1998)'s approach is similar to mine but in a one-shot context. The idea is that the null contract is attractive under max-min preferences because it gives each party a fixed fraction of the future surplus. This anticipates the key insight of

[3]My implementation result via the renegotiation of dynamically incomplete contracts is distinct from similar sounding results showing how optimal renegotiation-proof contracts in the expected utility setting are implementable with short-term contracts. See, for example, Fudenberg, Holmstrom, and Milgrom (1990). In that literature, short-term implementability is way to interpret the fact that the optimal contract is recursive over continuation payoffs and the value function is everywhere downward sloping. At any moment in time, in order to compute the current short-term contract, one must know the forward looking value function which effectively means that one has to also compute all possible short-term contracts that can occur in the future. In contrast, my implementation result implies that the parties can literally construct the optimal dynamic contract one incomplete step at a time. See Section 3.3 for a detailed comparison.

Carroll (2015), which looks at a static principal-agent model with moral hazard and max-min preferences, and shows that the optimal contract is linear. There are also a number other recent static contracting papers using max-min preferences to deliver simple, intuitive contracts in various settings. See, for example, Garrett (2012) and Frankel (2014).

Lastly, my analysis of debt under asymmetric information links two strands of the security design literature. First, there is the complete contracting approach that aims to show when the wedge payoff structure is optimal. See, for example, Townsend (1979), Gale-Hellwig (1985), DeMarzo and Duffie (1999), and Yang (2016). Second, there is the incomplete contracts approach that takes a particular contractual incompleteness as the primitive and is interested in the optimality of the state-contingent allocation of control rights induced by debt. See, for example, Aghion and Bolton (1992). I take the complete contracts approach but by using preference-for-robustness, I am also able to address issues involving renegotiation and contingent bargaining power allocation.

The rest of the paper is organized as follows. I begin with a toy model in Section 2 that introduces many of the basic ideas, terminology, and results. I then introduce the formal model in Section 3 and re-establish the results in the broader setting that will serve as the foundation for my analysis of contracting under asymmetric information. Finally, in Section 4, I introduce asymmetric information into the formal model and discuss the optimality debt.

2 A Toy Model

Suppose there are two players, call them E and F, who share a project that spans dates 0, 1, and 2. At date 0, the project has an initial investment k_0 of capital. For my purposes, it doesn't matter who provided the capital. The project will turn this investment into some new amount of capital v_1 at date 1. The players believe that the project is going to have an expected return of r_1 at date 1. More specifically, they believe $v_1 \sim U[0, 2k_0(1 + r_1)]$.

Once date 1 arrives, a specific value for v_1 is realized. In addition, an expected return r_2 linking investment at date 1 to production at date 2 is realized as well. At this point, the players can withdraw some of the project's capital for date 1 consumption subject to the obvious budget constraint: $c_{1,E}(v_1, r_2) + c_{1,F}(v_1, r_2) \leq v_1$. Here, I emphasize that the consumption amounts may depend on the date 1 state of the world which is summarized by v_1 and r_2. After consumption occurs, whatever capital is remaining in the project, call it k_1, then generates an expected return of r_2 at date 2. Again, the players believe that date 2 capital $v_2 \sim U[0, 2k_1(1 + r_2)]$. Finally at date 2, the players consume all of the realized v_2: $c_{2,E}(v_1, r_2, v_2) + c_{2,F}(v_1, r_2, v_2) = v_2$.

A contract specifies a pair of history-dependent consumption streams for E and F: $\{c_{1,E}(v_1, r_2), c_{1,F}(v_1, r_2), c_{2,E}(v_1, r_2, v_2), c_{2,F}(v_1, r_2, v_2)\}$ satisfying the aforementioned budget constraints. Given a contract, E's total payoff is $\mathbf{E}_{v_1,r_2}\left[u_E(c_{1,E}(v_1, r_2)) + \mathbf{E}_{v_2} u_E(c_{2,E}(v_1, r_2, v_2))\right]$ where u_E is some weakly concave strictly increasing utility function. F's total payoff is defined similarly. The optimal contracting problem is to characterize all renegotiation-proof contracts.

Solving the optimal contracting problem is standard. The solution uses backwards induction and is based on the fact that the state variable is the continuation payoff. In general, there is not much to say about the solution. There are some risk-sharing constraints that need to be satisfied. But the contract does not look like a dynamically incomplete contract being renegotiated over time. In fact, experience tells us that as one adds more frictions and features to the model, unless it is done in a delicate way, the solution to the optimal contracting problem will look more complex, and less like an incomplete contract.

What I want to show now is that with a few simple changes to the way E and F think about the future, I can turn the rather sensitive and complex optimal contract into something that looks like what was described in the beginning of the introduction. There are many ways to go about this and I will start with the most basic.

A. Incomplete Contracts with Ex-Post Renegotiation

As an easy first step, let's just focus on the players' belief about r_2 standing at date 0. Instead of having a specific distribution in mind, let's suppose that the players have no idea what r_2 will be and seek a renegotiation-proof contract will do well no matter what. That is, E (and similarly, F) evaluates long-term contracts as follows:

$$\min_{1+r_2\geq 0} \mathbf{E}_{v_1}\left[u_E(c_{1,E}(v_1,r_2)) + \mathbf{E}_{v_2} u_E(c_{2,E}(v_1,r_2,v_2))\right].$$

Notice, the expression is very similar to before, except there is now a min operator in front, capturing aversion to ex-ante fuzziness about r_2. Let us refer to this situation as one where E and F have "preference-for-robustness."[4]

When E and F have preference-for-robustness, any renegotiation-proof contract can be constructed in a simple, three-step process.

1. *Think about the "worst case" scenario.*

   Technically speaking, the worst case scenario for a player depends on the contract. However, there is an intuitive "worst case" scenario here, which is the scenario that $1+r_2=0$. So first, suppose the players think $1+r_2$ will equal zero for certain. In this case, any Pareto-optimal contract is just a static contract ending at date 1 that splits v_1 between E and F for date 1 consumption. Let $\alpha_E^*(v_1)+\alpha_F^*(v_1)=v_1$ denote this split.

2. *Make sure the contract does well in the "worst case" scenario.*

   Now, move back to the actual model, and introduce the long-term contract where at date 1, no matter the state of the world, $c_{1,E}(v_1,r_2)=\alpha_E^*(v_1)$, $c_{1,F}(v_1,r_2)=\alpha_F^*(v_1)$, and, consequently, $c_{2,E}\equiv c_{2,F}\equiv 0$.

3. *Make sure the contract still does well in other scenarios.*

   The contract right now is not renegotiation-proof because, depending on the realized r_2, it may in fact be better for both parties to save some capital at date 1. So, Pareto-improve each date 1 continuation contract to some individually rational Pareto-optimal

[4] At this point, preference-for-robustness looks like nothing more than just a special case of static ambiguity aversion. The dynamic max-min decision making aspect of preference-for-robustness will emerge once I further develop the concept.

continuation contract. Here, individual rationality is defined in relation to the prior split $(\alpha_E^*(v_1), \alpha_F^*(v_1))$.

The final product is a renegotiation-proof contract. The proof basically involves convincing oneself that having each date 1 Pareto-optimal continuation contract be individually rational in relation to $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ ensures that the date 0 renegotiation-proof constraint is also satisfied.

Notice, in the last step of constructing our renegotiation-proof contract, I do not specify how the parties must Pareto-improve $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ beyond the basic requirement that the end result must be Pareto-optimal. In fact, I don't need to, in the sense that, no matter which individually rational Pareto-optimal continuation contract is selected, the resulting renegotiation-proof contract delivers the same ex-ante payoff to both parties.

This is significant because it means that the players can evaluate a split $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ - in particular, decide whether or not it is optimal - and still be agnostic about how this preliminary agreement will be renegotiated if such a renegotiation is called for ex-post. Thus,

Result 1. *If players have preference-for-robustness then every point on the Pareto-frontier can alternatively be implemented by an incomplete contract with ex-post renegotiation: The players sign a date 0 incomplete contract specifying a split $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ which will serve as the outside option in date 1. This contract is incomplete because it does not specify how the players should Pareto-improve the outside option if possible. Instead, at date 1, if the parties find that they can do better than the outside option, then they will bargain at that time about which Pareto-optimal alternative to take.*

The little bit of methodology introduced so far can already be adapted to provide a simple rationale for equity:

B. Linear Incomplete Contracts

Begin with the preference-for-robustness setup described above. At this point, while the players have a fuzzy idea about r_2, they are absolutely clear that v_1 and v_2 are distributed according to a specific uniform distribution. This disparity in beliefs is jarring and unrealistic. So, let's now make the players have a fuzzy idea about how v_1 and v_2 are distributed as well.

At date 1, given a realized r_2, I continue to assume that the players believe capital invested will have an expected return of r_2 at date 2. However, they are no longer convinced that the distribution of v_2 must be $U[0, 2k_1(1 + r_2)]$. Instead, the players entertain any distribution of v_2 so long as the expected return is r_2. Similarly, at date 0, the players entertain any distribution of v_1 so long as the expected return is r_1. This means E (and similarly, F) evaluates long-term contracts as follows:

$$\min_{1+r_2\geq 0,\ \mathbf{E}_{\pi_1} v_1 = k_0(1+r_1)} \mathbf{E}_{\pi_1}\left[u_E(c_{1,E}(v_1, r_2)) + \min_{\mathbf{E}_{\pi_2} v_2 = k_1(1+r_2)} \mathbf{E}_{\pi_2} u_E(c_{2,E}(v_1, r_2, v_2))\right].$$

Call this new variant "expected value preference-for-robustness."

Under expected value preference-for-robustness, Result 1 holds. However, I can now say something more specific about the split $(\alpha^*_E(v_1), \alpha^*_F(v_1))$ and any potential renegotiation of the split.

Result 2. *If players have expected value preference-for-robustness then any Pareto-optimal allocation admits an implementation in linear incomplete contracts. Specifically, one can assume that the split $(\alpha^*_E(v_1), \alpha^*_F(v_1))$ is linear in v_1 and that any ex-post renegotiated continuation contract is linear in v_2.*

In this implementation of the optimal contract, the evolving linear split of v_t over time draws obvious comparisons to equity. What's interesting about this result is that equity is not some static arrangement that mechanically gives equity holders their permanently fixed shares of total consumption (i.e. dividends) over time. In fact, E and F's equity positions can vary over time and, moreover, the *actual* consumption amounts $(c_{t,E}, c_{t,F})$ at each date t can be highly nonlinear functions of v_t. This is because the linear splits interpreted as equity serve to only set the table for a dynamic renegotiation process which then ultimately determines consumption.

For example, the players at date 0 may begin with a linear split $(\alpha^*_E(v_1) = \alpha^*_E \cdot v_1, \alpha^*_F(v_1) = \alpha^*_F \cdot v_1)$ of v_1 but this does not mean they will literally consume $c_{1,E} = \alpha^*_E \cdot v_1$ and $c_{1,F} = \alpha^*_F \cdot v_1$ at date 1. Instead, given a date 1 state of the world s_1, the split will likely be renegotiated to some other continuation contract $(c^*_{1,E}(s_1), c^*_{1,F}(s_1), c^*_{2,E}(s_1, v_2), c^*_{2,F}(s_1, v_2))$. And while Result 2 says that one can always take $(c^*_{2,E}(s_1, v_2), c^*_{2,F}(s_1, v_2))$ to be a linear split of v_2, it does not say that this split must be $(\alpha^*_E(v_2), \alpha^*_F(v_2))$. Moreover, it does not say anything about $c^*_{1,E}(s_1)$ and $c^*_{1,F}(s_1)$. In particular, neither $c^*_{1,E}(s_1)$ nor $c^*_{1,F}(s_1)$ must linearly vary with v_1.

Proof. The proof is most easily understood by studying a classic static risk-sharing problem. Suppose E and F must agree on some sharing rule $(\alpha_E(v), \alpha_F(v))$ for some asset with a random one-time payoff v. If the players know exactly the distribution of v, then Borch's Law implies that any efficient sharing rule must equate the ratio of marginal utilities across v. That is, there exists some λ^* such that

$$\lambda^* = \frac{u'_E(\alpha_E(v))}{u'_F(\alpha_F(v))} \text{ for all } v.$$

Unless both parties have the same constant relative risk aversion, linear sharing is not efficient.

In contrast, suppose, instead, the players entertain an entire set $\mathcal{S}_{\overline{v}}$ of probability distributions for v where $\mathcal{S}_{\overline{v}}$ consists of all distributions π such that $\mathbf{E}_\pi v = \overline{v}$ for some constant $\overline{v}$. Suppose the players are averse to this fuzziness so that E (and similarly F) evaluates his share to be worth $\min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v)$. I now show that any efficient payoff can be achieved by a linear sharing rule.

Consider a sharing rule $(\alpha_E(v), \alpha_F(v))$ where $\alpha_E(v)$ is concave over, say, the region $[0, 1]$. Introduce an alternative sharing rule $(\tilde{\alpha}_E(v), \tilde{\alpha}_F(v))$ such that $\tilde{\alpha}_E(v)$ is identical to $\alpha_E(v)$ except on $[0, 1]$, where $\tilde{\alpha}_E(v) = (1 - v)\alpha_E(0) + v\alpha_E(1)$.

Obviously, F weakly prefers $(\tilde{\alpha}_E(v), \tilde{\alpha}_F(v))$ to $(\alpha_E(v), \alpha_F(v))$ since, by construction, he is getting weakly more for every v under $(\tilde{\alpha}_E(v), \tilde{\alpha}_F(v))$. The key step is to realize that E is indifferent between the two shares. Once this is proved, an immediate implication is that any efficient payoff can be achieved by a sharing rule where E's share is weakly convex. Then apply the logic once more, this time on F. Now the conclusion is that any efficient payoff can be achieved by a sharing rule where both players have a weakly convex share. The only sharing rules where both shares are weakly convex are precisely the linear sharing rules.

To see why E is indifferent between $(\tilde{\alpha}_E(v), \tilde{\alpha}_F(v))$ and $(\alpha_E(v), \alpha_F(v))$, fix an arbitrary distribution $\pi \in \mathcal{S}_{\overline{v}}$. Now consider another distribution $\tilde{\pi}$ which is identical to π except that all weight that was previously on $(0, 1)$ is now shifted to the end points $\{0, 1\}$ in a way so that the expected value of v is the same. Then $\tilde{\pi} \in \mathcal{S}_{\overline{v}}$. Due to the concavity of $\alpha_E(v)$ over $[0, 1]$, the expected payoff of $\alpha_E(v)$ to E declines moving from π to $\tilde{\pi}$. This implies $\min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v) \geq \min_{\tilde{\pi} \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v)$. But distributions of the form $\tilde{\pi}$ comprise a strict subset of $\mathcal{S}_{\overline{v}}$, so the reverse inequality is true as well. Thus, $\min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v) = \min_{\tilde{\pi} \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v)$.

Similarly, $\min_{\tilde{\pi} \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \tilde{\alpha}_E(v) = \min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \tilde{\alpha}_E(v)$. Finally, by construction, $\alpha_E(v)$ and $\tilde{\alpha}_E(v)$ have the same payoff under $\tilde{\pi}$. This implies $\min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \alpha_E(v) = \min_{\pi \in \mathcal{S}_{\overline{v}}} \mathbf{E}_\pi \tilde{\alpha}_E(v)$. □

So far, all of the variations of preference-for-robustness assume that the further out one looks, the fuzzier things are. But this is not always a realistic assumption. In many situations, one can imagine having an idea about how things might turn out a year from now and yet still not be able to elucidate how exactly one will get there day-by-day starting today. Let's try to produce a variation of preference-for-robustness that captures this situation:

C. Incomplete Contracts with Ex-Interim Renegotiation

Before introducing the new variant of preference-for-robustness, I need to enrich the model slightly. I now let $r_2 = r_2(k_1)$. This means that at date 1, if the players choose to take some capital away from the project, it can affect the returns to capital at date 2. In general, the more functional forms for $r_2(k_1)$ I allow the players to entertain, the easier it will be for me to get the contracting result I'm going after. However, any set of possible functions that includes all weakly increasing functions is more than sufficient. So I assume this.

At this point, let me remind the reader the sequence of events in the model. At date 0, capital k_0 is invested. At date 1, a capital amount v_1 and a return function $r_2(k_1)$ are realized. Next, the players consume $c_{1,E}(v_1, r_2(k_1)) + c_{1,F}(v_1, r_2(k_1)) \leq v_1$. The remaining unconsumed capital, k_1, is then reinvested with expected return $r_2(k_1)$. Finally, at date 2, capital v_2 is realized, which is fully consumed by the players, $c_{2,E}(v_1, r_2(k_1), v_2) + c_{2,F}(v_1, r_2(k_1), v_2) = v_2$.

Now I am ready to define the new "long horizon" preference-for-robustness. At date 1, the players are willing to entertain any distribution of v_2 so long as the expected return on capital is $r_2(k_1)$. Moving back to date 0, fix a belief π_1 of the date 1 state of the world $(v_1, r_2(k_1))$. Consider the quantity $\mathbf{E}_{\pi_1} v_1(1 + r_2(v_1))/k_0 - 1$. It is the expected return on capital from date 0 to date 2 given belief π_1, assuming that the project is left alone at date 1. Fix a constant ${}_0r_2$ and assume that the players at date 0 entertain any belief π_1 such that $\mathbf{E}_{\pi_1} v_1(1 + r_2(v_1))/k_0 - 1 = {}_0r_2$.

Here, I have just described players who, standing at date 0, know that the expected return to capital is ${}_0r_2$ from now to date 2. However, they have no idea how the project will evolve between now and date 2. For example, each day capital could grow at the constant rate $\sqrt{1+{}_0r_2}-1$, it could grow quickly to date 1 then slow down, or grow slowly to date 1 then speed up. Or, maybe there is a 50% chance the two date growth rate is well below ${}_0r_2$ and a 50% chance that it is well above ${}_0r_2$. The players entertain all of these scenarios and many more. With long horizon preference-for-robustness, E (and similarly, F) evaluates long-term contracts as follows:

$$\min_{\mathbf{E}_{\pi_1} v_1(1+r_2(v_1))/k_0 - 1 = {}_0r_2} \mathbf{E}_{\pi_1}\left[u_E(c_{1,E}(v_1, r_2(\cdot))) + \min_{\mathbf{E}_{\pi_2} v_2 = k_1(1+r_2(k_1))} \mathbf{E}_{\pi_2} u_E(c_{2,E}(v_1, r_2(\cdot), v_2))\right].$$

Under long horizon preference-for-robustness, solving the optimal contracting problem can be done using the same three step process described earlier, except, of course, the "worst case" scenario is now different (and a little less obvious). Consequently, Pareto-optimal allocations can be implemented in a way that is similar to the implementation described in Result 1, except the initial incomplete contract spans two dates and the renegotiation occurs ex-interim.

Result 3. *If players have long horizon preference-for-robustness then every point on the Pareto-frontier can alternatively be implemented by an incomplete contract with interim renegotiation: The players sign a date 0 incomplete contract that specifies a split* $(\alpha_E^*(v_2), \alpha_F^*(v_2))$ *of* v_2 *and leaves the project alone until date 2. At date 1, the continuation of this incomplete contract serves as the outside option. This contract is incomplete because it does not specify how the players should Pareto-improve the outside option if possible. Instead, at the interim date 1, if the parties find that they can do better than the outside option, then they will bargain at that time about which Pareto-optimal alternative to take.*

I will prove this result in the next section when I consider the formal model.

With Result 3 established, my analysis of the toy model is finished. This analysis has highlighted two important facets of preference-for-robustness. First, it showed how preference-for-robustness provides a micro-foundation for the type of dynamically incomplete contracts described in the introduction. The formal analysis in the next section will do this at a more general level but the main ideas are already there in this toy model section.

Second, it showed that the preference-for-robustness concept is flexible. Starting with a simple baseline notion of preference-for-robustness, I showed how the concept can be further specialized, leading to an implementation with sharper predictions (Result 2). I also showed how it can be generalized, leading to an implementation with richer features (Result 3). Looking ahead to Section 4, I will also show how the concept can be fruitfully "integrated" with the rest of the contracting literature by embedding a notion of preference-for-robustness into a general security design model with asymmetric information. As I described in the introduction, from that model will emerge a rich theory of debt. I now turn to the formal model.

3 Model

Players. There are two players, call them E for "entrepreneur" and F for "financier."

Project. There is a project v that spans dates 0, 1, and 2. At date 0, it consists of some amount of capital $v_0 \geq 0$. The players consume $c_{0,E} + c_{0,F} \leq v_0$. The remaining capital, call it k_0, then randomly generates an amount of date 1 capital $v_1 \geq 0$. Again the players consume some $c_{1,E} + c_{1,F} \leq v_1$ leading to k_1. Finally, k_1 randomly generates an amount of date 2 capital v_2 which is complete consumed, $c_{2,E} + c_{2,F} \leq v_2$.

States of the World. A date 2 state of the world s_2 is a realized date 2 capital v_2. Define $\{s_2\}$ to be the set of all possible date 2 states of the world. A date 1 state of the world s_1 is the following object,

$$(v_1, \Pi_2 : k_1 \to 2^{\Delta(\{s_2\})}). \tag{1}$$

A date 1 state of the world consists of a realized date 1 capital plus a *belief function* Π_2, which specifies for each date 1 leftover capital k_1 a set of beliefs about the date 2 state of the world. I impose some mild restrictions on the shape of Π_2. Specifically, I assume for every pair $k_1 < k_1'$, if $\pi_2 \in \Pi_2(k_1)$ then there exists a $\pi_2' \geq_d \pi_2$ that $\in \Pi_2(k_1')$, and if $\pi_2' \in \Pi_2(k_1')$ then there exists a $\pi_2 \leq_d \pi_2'$ that $\in \Pi_2(k_1)$. The partial order $\geq_d$ is by first-order stochastic dominance. I also assume $\Pi_2(0) \equiv 0$. Let $\{s_1\}$ denote the set of all possible date 1 states of the world. Finally, fix a subset $\mathcal{S}_1 \subset \{s_1\}$. Given $\mathcal{S}_1$, a date 0 state of the world s_0 is the following object,

$$(v_0; \Pi_1 : k_0 \to 2^{\Delta(\mathcal{S}_1)}). \tag{2}$$

It is defined in a similar way to the date 1 state of the world s_1. I impose the same mild restrictions on the shape of Π_1 as I did on the shape of Π_2. Also, notice, for every k_0, the support of every belief in $\Pi_1(k_0)$ is required to be a subset of $\mathcal{S}_1$.

Definition. *A setting is choice of* $(s_0, \mathcal{S}_1)$.

Contracts and Preferences. Fix a setting setting $(s_0, \mathcal{S}_1)$. A contract is a history dependent pair of consumption streams for E and F,

$$(c_E, c_F) := (c_{0,E}, c_{0,F}, c_{1,E}(s_1), c_{1,F}(s_1), c_{2,E}(s_1, s_2), c_{2,F}(s_1, s_2)).$$

A contract implies a sequence of leftover capital amounts $k := (k_0, k_1(s_1))$.

Given a contract (c_E, c_F). Then E's continuation payoff process U_E is defined as follows.

$$\begin{aligned} U_{2,E}(s_1, s_2) &= u_E(c_{2,E}(s_1, s_2)), \\ U_{1,E}(s_1) &= u_E(c_{1,E}(s_1)) + \min_{\pi_2 \in \Pi_2(k_1(s_1))} \mathbf{E}_{\pi_2} U_{2,E}(s_1, s_2), \\ U_{0,E} &= u_E(c_{0,E}) + \min_{\pi_1 \in \Pi_1(k_0)} \mathbf{E}_{\pi_1} U_{1,E}(s_1). \end{aligned}$$

Here u_E is E's strictly increasing, weakly concave utility function. u_F and U_F are defined similarly for F.

Lemma 1. *Preferences are weakly time-consistent.*

Proof. Without loss of generality, I prove the result for E. Fix a stopping time $\tau \in \{1, 2\}$, a sequence of leftover capital amounts k, and two consumption streams c' and c'' for E. Suppose that $c'_t = c''_t$ for all $t < \tau$ and $U'_{\tau,E}(s_1, \ldots, s_\tau) \leq U''_{\tau,E}(s_1, \ldots, s_\tau)$. Here U'_E and U''_E denote the continuation payoff processes from receiving c' and c'', respectively. Fix an s_1 such that $\tau(s_1) > 1$. Then $U'_{1,E}(s_1) = u_E(c'_1(s_1)) + \min_{\pi_2 \in \Pi_2(k_1(s_1))} \mathbf{E}_{\pi_2} U'_2(s_1, \cdot) = u_E(c''_1(s_1)) + \min_{\pi_2 \in \Pi_2(k_1(s_1))} \mathbf{E}_{\pi_2} U'_{2,E}(s_1, \cdot) \leq u_E(c''_1(s_1)) + \min_{\pi_2 \in \Pi_2(k_1(s_1))} \mathbf{E}_{\pi_2} U''_{2,E}(s_1, \cdot) = U''_{1,E}(s_1)$. In addition, for all s_1 such that $\tau(s_1) = 1$, it is assumed that $U'_{1,E}(s_1) \leq U''_{1,E}(s_1)$. Thus, for every s_1, $U'_{1,E}(s_1) \leq U''_{1,E}(s_1)$. Now, an argument similar to the one used above shows that $U'_{0,E} \leq U''_{0,E}$. □

The class of preferences I have introduced is quite broad. By varying the setting $(s_0, \mathcal{S}_1)$, it includes, as special cases, expected utility and all the variants of preference-for-robustness that I will consider in this paper. For now, here are two example settings:

Expected Utility Setting. $\mathcal{S}_1$ is the set of all date 1 states with a singleton valued Π_2. s_0 is a date 0 state with a singleton valued Π_1.

Preference-for-Robustness Setting. $\mathcal{S}_1 = \{s_1\}$. s_0 is a date 0 state of the world satisfying the following property: For each $k_0 > 0$, there is a set of distributions $V_1(k_0)$ of v_1 such that $\pi_1 \in \Pi_1(k_0)$ if and only if $\pi_1|_{v_1} \in V_1(k_0)$.

Just as in the toy model, I focus on renegotiation-proof contracts. So, from now on, the Pareto-frontier is the payoff frontier generated by renegotiation-proof contracts.

Theorem 1. *In the preference-for-robustness setting, fix a Pareto-optimal allocation $(U^*_{0,E}, U^*_{0,F})$. Then there exists a date 0 consumption pair $(c^*_{0,E}, c^*_{0,F})$ plus a split $\alpha^*_E(v_1) + \alpha^*_F(v_1) = v_1$ such that any contract with the same date 0 consumption and the property that for every state s_1, the continuation contract is Pareto-optimal with*

$$(U_{1,E}(h_1), U_{1,F}(h_1)) \geq (u_E(\alpha^*_E(v_1)), u_F(\alpha^*_F(v_1)))$$

is renegotiation-proof and achieves $(U^*_{0,E}, U^*_{0,F})$.

Proof. See Appendix. □

The theorem and its proof correspond to the three-step process for constructing renegotiation-proof contracts in the preference-for-robustness setting introduced in the toy model. Also just like in the toy model, renegotiation-proof contracts can be implemented by incomplete contracts with ex-post renegotiation.

3.1 Implementation with Incomplete Contracts.

Given a date 1 state s_1 and an outside option $(O_{1,E}, O_{1,F}) \in \mathbb{R}^2$ for E and F, a *date 1 renegotiation protocol* selects a Pareto-optimal continuation contract that is individually rational, or, if none exist, the outside option. A *date 0 incomplete contract* specifies a date 0 consumption plan plus, to serve as the outside option at date 1, a split of the project's future capital v_1.

In general, players cannot properly evaluate a date 0 incomplete contract without also specifying a date 1 renegotiation protocol. However, in the preference-for-robustness setting, players can evaluate a date 0 incomplete contract - in particular decide if it is optimal - without specifying beforehand how they would go about renegotiating it the next date if a renegotiation is warranted. Given this fact, it is well-defined to say,

Corollary 1. *Under preference-for-robustness, any Pareto-optimal allocation can be implemented by an incomplete contract with ex-post renegotiation.*

Proof. Let the date 0 incomplete contract be $(c^*_{0,E}, c^*_{0,F}, \alpha^*_E(v_1), \alpha^*_F(v_1))$. □

This corollary corresponds to Result 1 in the toy model. Following the toy model, next I introduce a variant of preference-for-robustness and derive a linearity result.

Expected Value Preference-For-Robustness Setting. $\mathcal{S}_1$ is the set of all date 1 states with a singleton valued Π_2. s_0 is a date 0 state with a singleton valued Π_1.

Corollary 2. *[Linear Incomplete Contracts] In the expected value preference-for-robustness setting, every Pareto-optimal allocation can be achieved by a contract where* $(\alpha^*_E(v_1), \alpha^*_F(v_1))$ *is a linear split of* v_1*, and each date 1 continuation contract's date 2 consumption plan is a linear split of* v_2*.*

This corollary corresponds to Result 2 in the toy model.

The reduction from an arbitrary incomplete contract to a linear one is mathematically similar to the ones appearing in Mukerji (1998) and Carroll (2015). Indeed, the example can be seen as a way of extending their linearity result to a dynamic context. Two notable differences are that in my model neither party is required to be risk-neutral and the result holds even allowing for random contracts.

A natural interpretation of this linear incomplete contract is as an equity contract, where both parties start with some initial equity position (α^*_E, α^*_F) and can rebalance over time based on changing risk-attitudes and opinions about the future viability of the project. Recall the discussion in Section 2.

3.2 Beyond The Formal Model

The paper focuses on a three date model purely for notational simplicity.The formal model and all results generalize easily to a setting with $T + 1$ dates for $T > 2$. At date T, a state world s_T is a realized v_T. For all $t < T$, the state of the world s_t is defined by backwards induction: Let $\{s_{t+1}\}$ denote the set of all possible date $t + 1$ states of the world. Fix an $\mathcal{S}_{t+1} \subset \{s_{t+1}\}$. Then a date t state of the world is

$$s_t = (v_t, \Pi_{t+1} : k_t \to 2^{\Delta(\mathcal{S}_{t+1})}).$$

A setting is a choice $(s_0, \mathcal{S}_1, \dots, \mathcal{S}_{T-1})$. Just like before, different settings will produce expected utility and all variants of preference-for-robustness.

Under preference-for-robustness, Pareto-optimal contracts can be implemented with incomplete contracts with successive ex-post renegotiation, generalizing Theorem 1 and Corollary 1. At date 0, the players agree to a date 0 consumption plan plus a split $(\alpha^*_{1,E}, \alpha^*_{1,F})$ of v_1. At date 1, the realized split $(\alpha^*_{1,E}(v_1), \alpha^*_{1,F}(v_1))$ gets renegotiated to a new date 1 consumption plan plus a split $(\alpha^*_{2,E}, \alpha^*_{2,F})$ of v_2. In general, at date $t < T$, the realized split $(\alpha^*_{t,E}(v_t), \alpha^*_{t,F}(v_t))$ gets renegotiated to a new date t consumption plan plus a split $(\alpha^*_{t+1,E}, \alpha^*_{t+1,F})$ of v_{t+1}.

The formal model can also be generalized to allow for players to have differing opinions about the future. For example, take the three date formal model. Keep the date 2 state of the world the same. Now define a date 1 state of the world s_1 to be

$$(v_1, \Pi_{2,E}, \Pi_{2,F} : k_1 \to 2^{\Delta(\{s_2\})}).$$

Again, fix a subset $\mathcal{S}_1 \in \{s_1\}$ and now a date 0 state of the world s_0 is

$$(v_0, \Pi_{1,E}, \Pi_{1,F} : k_0 \to 2^{\Delta(\mathcal{S}_1)}).$$

All variants of preference-for-robustness admit natural generalizations to allow for potentially diverging belief functions over time. Theorem 1, Corollary 1, and Corollary 2 continue to hold.

3.3 Comparisons to Contracting with Bayesian Players

How is the implementation results different from those in the Bayesian contracting literature about renegotiation-proof optimal long-term arrangements being implementable as sequences of short-term contracts?

Consider a contracting model where Pareto-optimal long-term contracts are recursive over continuation payoffs. Moreover, assume that the continuation payoff process always lies on the Pareto-frontier. Then there is a way to describe each Pareto-optimal long-term contract so that it is as if one party is contracting the other only through a series of short-term contracts. For example, Pareto-optimal contracts in the Bayesian version of my model satisfy the above assumptions. As a result, Pareto-optimal contracts can be described as follows: At each date t, given a state of the world s_t and a continuation payoff $U_{t,E}$ promised to E, F gives E a short-term contract specifying a $c_{t,E}(s_t, U_{t,E})$, $k_t(s_t, U_{t,E})$, and a state-contingent promised continuation payoff $U_{t+1,E}(s_{t+1}, s_t, U_{t,E})$ for the next date. Then, when the next date rolls around and s_{t+1} and $U_{t+1,E}$ are realized, F writes another short-term contract and so on.

This type of reinterpretation of optimal long-term arrangements has appeared many times in the Bayesian renegotiation-proof contracting literature.[5] Oftentimes this is interpreted to

[5]The renegotiation-proof qualifier is needed in order to guarantee that the continuation payoff always lies on the Pareto-frontier, which, recall, is a necessary condition for the reinterpretation. See, for example, Fudenberg, Holmstrom and Milgrom (1990).

mean that these arrangements, while derived in a complete contracts, Bayesian setting, can nevertheless be practically thought of as dynamically incomplete contracts being renegotiated over time.

How does this result contrast with Corollary 1, which interprets optimal contracts under preference-for-robustness as dynamically incomplete contracts being renegotiated over time, but is not applicable to Bayesian optimal contracts?

Imagine you are a lawyer hired by E and F. You have access to all relevant information and your job is to craft a Pareto-optimal contract for the two parties. Before you begin, E and F reassure you that it is not such a hard task by explaining to you that at each date, you only need to write a short-term contract that lasts until the next date.

Is this of any comfort to you?

Suppose E and F are Bayesian. In order to figure out the optimal state-contingent promised continuation payoff $U_{t+1,E}(s_{t+1}, s_t, U_{t,E})$, you need to know the date $t+1$ value function. In order to compute the date $t+1$ value function, you need to know for each possible state s_{t+1}, what short-term contract you would write, which then requires you to know the date $t+2$ value function and so on. Thus, to be able to compute just the single short-term contract for today, you need to compute all possible short-term contracts that can come after it at all dates in the future.

Now suppose E and F have preference-for-robustness. In order to compute today's incomplete contract, you do not need to compute any future incomplete contract. You do not need to consider E and F's beliefs in the future. To compute today's incomplete contract, all you need to think about is E and F's beliefs today about the liquidation value tomorrow. This allows you to construct the Pareto-optimal contract one incomplete step at a time. As a result, you end up only computing the short-term contracts for states that actually get realized rather than for all possible states that could have been realized.

How is having preference-for-robustness different from believing the world is going to end?

Another concern is how different are Theorem 1 and Corollary 1 from optimal contracts arising in settings where players are Bayesian but one or both players think "the world is going to end after tomorrow" (i.e. at date 0, the belief function only places weight on those states s_1 that have $\Pi_{2,\cdot} \equiv 0$).

For example, suppose both E and F are Bayesian, and only F thinks the world is going to end after date 1. Then the optimal contract will simply load all of F's date 1 consumption on those states s_1 with $\Pi_{1,F} \equiv 0$. In all other states, E gets the entire project.

This does not look at all like the optimal contract arising under preference-for-robustness.

Next, suppose E and F are Bayesian, and both think the world is going to end after date 1. Then over the set of date 1 states where $\Pi_{1,E}$ and/or $\Pi_{1,F} \equiv 0$, the optimal contract fully liquidates the project and splits the proceeds between E and F. For all other date 1 states, anything goes.

Here, the full liquidation of the project followed by the split of v_1 over those states where $\Pi_{1,E}$ and/or $\Pi_{1,F} \equiv 0$, is reminiscent of the feasible split $(\alpha^*_E(v_1), \alpha^*_F(v_1))$ in Theorem 1. However, there is nothing tying this split of v_1 to the continuation payoffs for E and F in those states where neither $\Pi_{1,E}$ nor $\Pi_{1,F}$ are trivial. This is unlike what happens under preference-for-robustness, where it is crucial that all continuation contracts deliver payoffs that dominate the feasible split $(\alpha^*_E(v_1), \alpha^*_F(v_1))$. The discipline imposed on continuation payoffs across a wide range of states under preference-for-robustness is what allows the optimal contract to be interpreted as a dynamically incomplete contract that is renegotiated over time. In the end-of-the-world model, one could take an optimal contract and change it so that neither party gets any consumption if the world doesn't end after date 1, and the contract would remain optimal. However, this contract is worth zero to both E and F under preference-for-robustness.

In the coming sections, I further develop the notion of preference-for-robustness and apply it to extensions of the model. Optimal contracts will be implemented with longer horizon incomplete contracts featuring interim renegotiation and refinance-able debt featuring state-contingent allocation of control rights. These results will help further distinguish optimal contracting under preference-for-robustness from expected utility optimal contracting.

3.4 Incomplete Contracts with Interim Renegotiation

So far I've shown how optimal complete contracts can be implemented as incomplete contracts featuring ex-post renegotiation. Empirical work by Roberts and Sufi (2009) and Roberts (2015) shows, however, that the typical loan contract is renegotiated early and frequently, before the initial contract matures. Moreover, much of this renegotiation is not triggered by the actual or anticipated violation of a pre-specified covenant. I now extend the notion of preference-for-robustness to provide micro-foundation for incomplete contracts with interim renegotiation. This analysis mirrors the last part of the toy model section.

Long Horizon Preference-for-Robustness Setting. $\mathcal{S}_1 = \{s_1\}$. s_0 is a date 0 state of the world satisfying the following property: Fix an arbitrary $\pi_1 \in \Delta(\mathcal{S}_1)$. Define $\Pi_2 \circ \pi_1 := \{\pi_1 \cdot \boldsymbol{\pi}_2 \mid \boldsymbol{\pi}_2(s_1) \in \Pi_2(v_1) \text{ for all } s_1\}$ to be the set of all probability distributions of v_2 generated by π_1 assuming the project is left alone at date 1. For each k_0, there is a set of probability distributions of v_2, call it $V_2(k_0)$, such that $\pi_1 \in \Pi_1(k_0)$ if and only if $\Pi_2 \circ \pi_1 \subset V_2(k_0)$.

Theorem 2. *In the long horizon preference-for-robustness setting, fix a Pareto-optimal allocation $(U^*_{0,E}, U^*_{0,F})$. Then there exists a date 0 consumption pair $(c^*_{0,E}, c^*_{0,F})$ plus a split $\alpha^*_E(v_2) + \alpha^*_F(v_2) = v_2$, such that any contract with the same date 0 consumption and the property that for every state s_1, the continuation contract is Pareto-optimal with*

$$(U_{1,E}(s_1), U_{1,F}(s_1)) \geq \left(\min_{\pi_2 \in \Pi_2(v_1)} \mathbf{E}_{\pi_2} u_E(\alpha^*_E(v_2)), \min_{\pi_2 \in \Pi_2(v_1)} \mathbf{E}_{\pi_2} u_F(\alpha^*_F(v_2)) \right). \tag{3}$$

is renegotiation-proof and achieves $(U^*_{0,E}, U^*_{0,F})$.

Fix a date 1 state s_1 and a continuation contract to serve as the outside option. An *date 1 interim renegotiation protocol* selects a Pareto-optimal continuation contract that is individually rational given the outside option.

A *date 0 incomplete contract* specifies a date 0 consumption plan, leaves the project alone at date 1, and specifies a feasible split of v_2. Fix a date 0 incomplete contract and, for each state s_1, let the continuation of the date 0 incomplete contract serve as the outside option at date 1. Then, as long as some date 1 interim renegotiation protocol is used, each player's valuation of the date 0 incomplete contract is invariant over the protocol choice. Just like before, players can determine when such a date 0 incomplete contract is optimal without specifying beforehand how they would go about renegotiating it the next date if such an interim renegotiation is warranted. Thus, it is well-defined to say,

Corollary 3. *Any Pareto-optimal allocation can be implemented by an incomplete contract with interim renegotiation.*

This corollary corresponds to Result 3 in the toy model. My analysis of the perfect information model is now complete. In the next and final section of the paper, I turn my attention to contracting under asymmetric information.

4 Debt and Control Rights

Standard debt has two salient features. First, the payoff function is a wedge function of the form $D \wedge v$ where D is a constant and v is the underlying asset value. Second, the allocation of control rights is done in a state-contingent way where the financier retains control of the asset if and only if the entrepreneur is in default. Implied by the allocation of control rights is the notion that the wedge payoff function is preliminary and may eventually be renegotiated.

The security design literature looking at debt largely splits into two groups, depending on which of these two features is the focal point. There is a complete contracts approach that focuses on explaining the wedge payoff function. In addition, there is an incomplete contracts approach, where a primary objective is to highlight how variations in the way securities allocate control rights can also have important payoff implications.

My goal is to unify these two approaches. I do this by introducing an asymmetric information problem into the formal model. The entrepreneur privately observes the true date 1 state of the world and can strategically reveal a contractible reported date 1 state of the world. I then show that the resulting optimal complete contract can be implemented by an incomplete contract featuring the wedge payoff function. This contract is then renegotiated according to a protocol that allocates bargaining power in a state-contingent way based on whether the project is in technical default (i.e. when the capital level is less than the debt obligation).

Before defining the asymmetric information problem, I first need to enrich the model. At date 1, I assume that a temporary taste shock $\theta_{1,E}$ is realized. Given $\theta_{1,E}$, E's utility for date 1 is $\theta_{1,E} u_E(\cdot)$. At date 2, it reverts back to $u_E(\cdot)$.[6]

States of the World. For simplicity, I assume asymmetric information exists only at date 1.

[6]The particular way the taste shock affects the utility function is not important. More generally, what I need is for there to be a possibility that E values consumption tomorrow much more than today. This can be achieved by a taste shock but it can also be achieved by allowing the players to have potentially diverging opinions about the future viability of the project and having E sometimes be very optimistic about the future. Recall Section 3.2 where I discuss how the model can be easily generalized to allow for differing belief functions for E and F.

Thus, the date 2 state of the world, which is a realized capital amount v_2, is still publicly observed. The date 1 state of the world s_1 is privately observed by E and is the following object,

$$(v_1, \theta_{1,E}, \Pi_2 : (k_1, \delta) \to 2^{\Delta(\{s_2\})}). \tag{4}$$

Notice, the "belief function" Π_2 is different compared to before. It is now a function of k_1 and an as yet undefined parameter δ. For now, I set Π_2 aside.

At date 1, upon observing s_1, E chooses to credibly report a portion of the capital $\hat{v}_1 \leq v_1$. The budget constraint is then $c_{1,E} + c_{1,F} \leq \hat{v}_1$. The reported leftover capital is $\hat{k}_1 := \hat{v}_1 - (c_{1,E} + c_{1,F})$ which is weakly smaller than the true value k_1 that is known only to E. In addition, E reports a taste shock $\hat{\theta}_{1,E}$.

Define $\delta := k_1 - \hat{k}_1 = v_1 - \hat{v}_1$ to be the amount of underreporting by E. Given reports $(\hat{v}_1, \hat{\theta}_{1,E})$, the induced public reported date 1 state of the world is

$$\hat{s}_1 := (\hat{v}_1, \hat{\theta}_{1,E}, \hat{\Pi}_1 : \hat{k}_1 \to 2^{\Delta(\{s_2\})}) \tag{5}$$

where $\hat{\Pi}_1(\cdot) := \Pi_1(\cdot + \delta, \delta)$. The first two parts of $\hat{s}_1$ I have already explained. $\hat{\Pi}_1$ is F's induced belief function given E's underreporting as measured by δ. It is defined over $\hat{k}_1$, which is what F believes is the amount of capital in the project at the end of date 1. As an abuse of notation, let $\Pi_1(k_1) := \Pi_1(k_1, 0)$. $\Pi_1(k_1)$ is F's belief function when E tells the truth. I assume that it is also E's belief function. Thus, if E tells the truth, then the model is identical to the perfect information model.

I impose two mild restrictions on the shape of $\Pi_1(k_1, \delta)$. For every δ, the induced belief function $\hat{\Pi}_1$ is weakly increasing in $\hat{k}_1$ with $\hat{\Pi}_1(0) = 0$ just like in the perfect information setting. Also, I assume that Π_1 is weakly decreasing in δ. This means the more E underreports v_1, the weakly more pessimistic is F.

Let $\{s_1\}$ denote the set of all date 1 states of the world. Given a subset $\mathcal{S}_1 \subset \{s_1\}$, the definition of s_0 is similar to before and is publicly observed.

A contract specifies a consumption plan (c_E, c_F) along with a report strategy $(\hat{v}_1, \hat{\theta}_{1,E})$ for E. A contract is truth-telling if $\hat{v}_1 = v_1$ and $\hat{\theta}_{1,E} = \theta_{1,E}$. As an abuse of notation, I let $s_1 = (v_1, \theta_{1,E}, \Pi_1 : k_1 \to 2^{\Delta(\{s_2\})})$ denote the reported state when the privately observed state is $s_1 = (v_1, \theta_{1,E}, \Pi_{1,F} : (k_1, \delta) \to 2^{\Delta(\{s_2\})})$ and E tells the truth.

Again, I focus on renegotiation-proof contracts. Since there is asymmetric information at date 1, defining the renegotiation-proof condition at date 1 requires care.

Fix a truth-telling contract with consumption plan (c_E, c_F) and some state s_1. Consider a (possibly off-equilibrium) reported state $\hat{s}_1$. The continuation contract is some

$$(c_{1,E}(\hat{s}_1), c_{1,F}(\hat{s}_1), c_{2,E}(\hat{s}_1, \cdot), c_{2,F}(\hat{s}_1, \cdot)).$$

A *renegotiation* of this continuation contract is some alternate consumption plan

$$(c'_{1,E}, c'_{1,F}, c'_{2,E}(\cdot), c'_{2,F}(\cdot))$$

satisfying the budget constraint $c'_{1,E} + c'_{1,F} \leq \hat{v}_1$. Under the renegotiation, the true leftover capital is $k'_1 = v_1 - (c'_{1,E} + c'_{1,F})$ and the reported leftover capital is $\hat{k}'_1 = \hat{v}_1 - (c'_{1,E} + c'_{1,F})$.

Under the renegotiation, E's continuation payoff is

$$u_E(c'_{1,E}) + \min_{\pi_2 \in \Pi_2(k'_1)} \mathbf{E}_{\pi_2} u_E(c'_{2,E}(\cdot)), \tag{6}$$

and F thinks his continuation payoff is

$$u_F(c'_{1,F}) + \min_{\pi_2 \in \hat{\Pi}_2(\hat{k}'_1)} \mathbf{E}_{\pi_2} u_F(c'_{2,F}(\cdot)). \tag{7}$$

E is strictly better off under the renegotiation if

$$(6) \geq u_E(c_{1,E}(s_1)) + \min_{\pi_2 \in \Pi_2(k_1)} \mathbf{E}_{\pi_2} u_E(c_{2,E}(s_1, \cdot)),$$

and F thinks he is strictly better off under the renegotiation if

$$(7) \geq u_F(c_{1,F}(\hat{s}_1)) + \min_{\pi_2 \in \hat{\Pi}_2(\hat{k}_1)} \mathbf{E}_{\pi_2} u_F(c_{2,F}(\hat{s}_1, \cdot)).$$

Definition. *A contract is renegotiation-proof at date 1 if and only if there does not exist a state s_1 and a renegotiation of the continuation contract following some reported state $\hat{s}_1$ that makes E strictly better off and makes F think it is strictly better off.*

Notice, if a truth-telling contract is renegotiation-proof at date 1, then truth-telling is automatically incentive-compatible.

I am almost ready to state the debt result. In the security design literature, it is common to impose some monotonicity constraint on the contracting space in order to derive a debt result. I could do that here as well, but I do not need to if, instead, I use the following variant of preference-for-robustness.

Lower Bound Preference-For-Robustness Setting. $\mathcal{S}_1 = \{s_1\}$. s_0 is a date 0 state of the world satisfying the following property: For each $k_0 > 0$, there is a set of distributions $V_1(k_0)$ of v_1 such that $\pi_1 \in \Pi_1(k_0)$ if and only if there is some $\mu \in V_1(k_0)$ such that $\pi_1|_{v_1} \geq_d \mu$.

Theorem 3. *In the lower bound preference-for-robustness setting, any Pareto-optimal allocation can be achieved by a Pareto-optimal contract with the following structure: There exists some constant $D_1 \geq 0$ such that for every date 1 state s_1, the continuation contract maximizes $U_{1,E}(s_1)$ subject to $U_{1,F}(s_1) = u_F(D_1)$ if $v_1 \geq D_1$, and maximizes $U_{1,F}(s_1)$ if $v_1 < D_1$.*

Proof. See Appendix. □

Definition. *A debt contract consists of the following two objects:*

1. *A date 0 incomplete contract where the feasible split of v_1 that serves as the date 1 outside option takes the form, $(\alpha^*_E(v_1), \alpha^*_F(v_1)) = (v_1 - v_1 \wedge D_1, v_1 \wedge D_1)$.*

2. *The date 1 renegotiation protocol that has E make a take-it-or-leave-it offer when* $v_1 \geq D_1$ *and has F make a take-it-or-leave-it offer when* $v_1 < D_1$.

Corollary 4. *Any Pareto-optimal allocation can be implemented by a debt contract.*

To get a feel for how the proof of Theorem 3 works, let us take the debt contract, keep the date 0 incomplete contract the same but change the renegotiation protocol so that E gets to make a take-it-or-leave-it offer all the time. I will now informally argue that this contract is no good for F:

First, it is important to remember that, in general, the best take-it-or-leave-it offer E can make that delivers payoff $D_1 \wedge v_1$ to F depends on the realized state of the world. In particular, if the realized belief function is optimistic about the future and E's realized taste shock takes a relatively low value, then it may be optimal for both parties to forgo consuming much at date 1, and, instead, wait for date 2's "bounty."

Let us assume such a state has occured at date 1. Moreover, assume that the belief function is relatively insensitive to δ, so that underreporting by E does not seriously adversely affect F's belief.

Then, I claim that E is better underreporting to a lower value $\hat{v}_1 < v_1$. Why?

Embedded in the structure of a debt contract is an obvious temptation to underreport - the less E reports, the less he has to deliver to F. However, there are some potential downsides to underreporting that may overcome this upside. For one thing, reporting a lower capital value tightens the date 1 budget constraint, which may be painful for E if he wants to consume a lot at date 1. Moreover, if underreporting makes F sufficiently pessimistic, then, even though E now owes F less, he may still have to give up strictly more to deliver that lesser continuation payoff due to the increased pessimism.

But notice, neither of these two potential downsides are a concern for us. We have assumed a state where E does not want to consume much at date 1 anyways and we have assumed that underreporting does not make F much more pessimistic. Thus, in this particular state, E can profitably underreport at the expense of F.

Now, of course, this is just one special state where underreporting is attractive. No doubt there are many other states with the same capital v_1 where E would not want to underreport. But remember, the setting is one of preference-for-robustness. Thus, even though the "worst case" scenario of underreporting does not always occur, it matters at the margin.

Our critical assessment of the altered debt contract suggests the following lesson,

Remark. *Varying F's payoff with project value when E is in control can backfire.*

The implication is that either F's payoff should be constant or if this is not desirable or feasible then F needs to be in control. A debt contract has precisely this property. The formal proof of Theorem 3 is basically a more fleshed out version of the arguments made here.

Just how bad can giving E control be when the project is in default?

Corollary 5. *Fix an arbitrary debt contract and change the renegotiation protocol to any one where E makes a take-it-or-leave-it offer when* $v_1 < D_1$. *Then the modified debt contract does not always induce truth-telling. Moreover, the payoff of the contract to F is zero.*

The implementation of the optimal contract by debt differs from the previous implementation results in one key way: The renegotiation protocol can no longer be aribtrary. As we have seen, the importance of allocating bargaining power in a state-contingent way emerges as an optimal reponse to the information asymmetry of the model.

Without this information asymmetry or, equivalently, if E could somehow be forced to tell the truth, then once again the debt contract's renegotiation protocol can be changed to any other and the ex-ante payoff of the contract to both parties would be unaffected. That is, the choice of the renegotiation protocol still does not have any direct effect on payoffs. It only indirectly affects payoffs through its affect on E's best-response reporting strategy. This means that the optimal long-term contract can still be constructed one incomplete step at a time like in the perfect information setting.

In related work, Townsend (1979) and Gale-Hellwig (1985) study security design under asymmetric information with costly state verification. They show in a standard expected-utility setting how debt might also emerge as the optimal contract. However, there are limitations to the result. Both parties must be risk-neutral and random contracts are disallowed. Moreover, the result does not extend to a dynamic setting and the contract is not renegotiation-proof. As is well-known, these assumptions are not mild. For example, Mookherjee and Png (1989) shows that if random state verifications are allowed then the optimal contract looks drastically different from debt. In contrast, my debt result is derived in a dynamic setting, focusing on renegotiation-proof contracts, where players can be risk-averse and contracts can be random.

Future Directions

In this paper I introduced a baseline notion of preference-for-robustness and then looked at a number of variations and their implications for optimal dynamic contracting. There are obviously many more natural variations that one can consider, and some of them may lead to surprising new insights. Another direction is to keep the preference-for-robustness as is and make the model richer. For example, one can introduce moral hazard or multilateral financing contracts. It would be interesting to see, for example, how the optimal debt contract in the asymmetric information model changes under these enrichments.

Appendix

Proof of Theorem 1. Fix a Pareto-optimal allocation $(U^*_{0,E}, U^*_{0,F})$ and a contract (c_E, c_F) that achieves it. For each v_1, there is a date 1 state s'_1 with capital v_1 such that $k_1(s'_1) = 0$. This will occur whenever the players are sufficiently pessimistic about the future.

For each v_1, define $\alpha^*_E(v_1) := c_{1,E}(s'_1)$ and $\alpha^*_F(v_1) := c_{1,F}(s'_1)$. For each $\pi_1 \in \Pi_1(k_0)$, define $\pi'_1 \in \Pi_1(k_0)$ to be a belief with the same marginal distribution over v_1 but has support only over those states of the form s'_1. Such a belief can be found in $\Pi_1(k_0)$ due to the preference-for-robustness assumption. Then

$$\begin{aligned}
U^*_{0,E} &= u_E(c_{0,E}) + \min_{\pi_1 \in \Pi_1(k_0)} \mathbf{E}_{\pi_1} U_{1,E}(s_1) \\
&\leq u_E(c_{0,E}) + \min_{\pi_1' \in \Pi_1(k_0)} \mathbf{E}_{\pi_1'} U_{1,E}(s_1') \\
&= u_E(c_{0,E}) + \min_{\pi_1' \in \Pi_1(k_0^*)} \mathbf{E}_{\pi_1'} u_E(\alpha^*_E(v_1)) \\
&= u_E(c_{0,E}) + \min_{\pi_1 \in \Pi_1(k_0^*)} \mathbf{E}_{\pi_1} u_E(\alpha^*_E(v_1))
\end{aligned}$$

Similarly, $U^*_{0,F} \leq u_F(c_{0,F}) + \min_{\pi_1 \in \Pi_1(k_0)} \mathbf{E}_{\pi_1} u_F(\alpha^*_F(v_1))$. But equality can be achieved for both continuation payoffs simply splitting v_1 according to $(\alpha^*_E(v_1), \alpha^*_F(v_1))$ no matter what is the date 1 state. This new contract achieves the Pareto-optimal allocation but is not renegotiation-proof. To make it renegotiation-proof, weakly Pareto-improve each continuation contract so that it becomes a Pareto-optimal continuation contract with $(U_{1,E}(s_1), U_{1,F}(s_1)) \geq (u_E(\alpha^*_E(v_1)), u_F(\alpha^*_F(v_1)))$. The resulting contract, call it (c^*_E, c^*_F), weakly increases each player's continuation payoff after every history s_1. Thus, (c^*_E, c^*_F) achieves the Pareto-optimal allocation $(U^*_{0,E}, U^*_{0,F})$. Moreover, since every continuation contract is Pareto-optimal by design, (c^*_E, c^*_F) is renegotiation-proof. □

Proof of Theorem 2. If either $U^*_{0,E}$ or $U^*_{0,F}$ is minimal then the theorem is obviously true. So assume neither payoff is minimal.??????

Fix a contract (c^*_E, c^*_F) that achieves $(U^*_{0,E}, U^*_{0,F})$ with date 0 leftover capital k^*_0. Fix an $\varepsilon > 0$ and define the date 1 state of the world

$$\hat{s}^\varepsilon_1 = \left(\varepsilon; \hat{\Pi}_2 : k_1 \to 2^{\Delta(\{s_2\})}\right)$$

with the following properties: $\hat{\Pi}_2(\varepsilon) = V_2(k^*_0)$, and $\hat{\Pi}_2(\delta) = 0$ for all $\delta < \varepsilon$. The belief $\hat{\pi}^\varepsilon_1 \in \Delta(\mathcal{S}_1)$ that puts all weight on $\hat{s}^\varepsilon_1$ is in $\Pi_1(k^*_0)$.

Consider the continuation contract following $\hat{s}^\varepsilon_1$. If there is any date 1 consumption by either player, then $(U_{1,E}(\hat{s}^\varepsilon_1), U_{1,F}(\hat{s}^\varepsilon_1)) \leq (u_E(\varepsilon), u_F(\varepsilon))$. This combined with the fact that $\hat{\pi}^\varepsilon_1 \in \Pi_1(k^*_0)$ means that if ε is sufficiently small, there will be no consumption at date 1 and $k_1(\hat{s}^\varepsilon_1) = \varepsilon$. In this case, define $\alpha^*_E(v_2) := c^*_{2,E}(\hat{s}^\varepsilon_1, v_2)$ and $\alpha^*_F(v_2) := c^*_{2,F}(\hat{s}^\varepsilon_1, v_2)$. Thus, the continuation payoffs satisfy

$$\begin{aligned}
U^*_{1,E}(\hat{s}^\varepsilon_1) &= \min_{\pi_2 \in V_2(k^*_0)} \mathbf{E}_{\pi_2} u_E(\alpha^*_E(v_2)), \\
U^*_{1,F}(\hat{s}^\varepsilon_1) &= \min_{\pi_2 \in V_2(k^*_0)} \mathbf{E}_{\pi_2} u_F(\alpha^*_F(v_2)).
\end{aligned}$$

Since $\hat{\pi}^\varepsilon_1$ is just one element of $\Pi_1(k^*_0)$, it must be that

$$\begin{aligned}
U^*_{0,E} &\leq u_E(c^*_{0,E}) + U^*_1(\hat{s}^\varepsilon_1), \\
U^*_{0,F} &\leq u_F(c^*_{0,F}) + U^*_1(\hat{s}^\varepsilon_1).
\end{aligned}$$

However, equality can be achieved by having the continuation contract after every s_1 be the same as the one after $\hat{s}_1^\varepsilon$. To see this, note that for every belief $\pi_1 \in \Pi_1(k_0^*)$, we have

$$\begin{aligned} U_{0,E} &= u_E(c_{0,E}^*) + \min_{\pi_1 \in \Pi_1(k_0^*)} \mathbf{E}_{\pi_1} U_{1,E}(s_1) \\ &= u_E(c_{0,E}^*) + \min_{\pi_1 \in \Pi_1(k_0^*)} \mathbf{E}_{\pi_1} \min_{\pi_2 \in \Pi_2(v_1)} \mathbf{E}_{\pi_2} u_E(\alpha_E^*(v_2)) \\ &= u_E(c_{0,E}^*) + \min_{\pi_1 \in \Pi_1(k_0^*)} \min_{\pi_2 \in \Pi_2 \circ \pi_1} \mathbf{E}_{\pi_2} u_E(\alpha_E^*(v_2)) \\ &\geq u_E(c_{0,E}^*) + U_{1,E}^*(\hat{s}_1^\varepsilon). \end{aligned}$$

A similar argument proves $U_{0,F}^* \geq u_F(c_{0,F}^*) + U_{1,F}^*(\hat{s}_1^\varepsilon)$.

Now the continuation payoff after each history h_1 is

$$\left(\min_{\pi_2 \in \Pi_2(v_1)} \mathbf{E}_{\pi_2} u_E(\alpha_E^*(v_2)), \min_{\pi_2 \in \Pi_2(v_1)} \mathbf{E}_{\pi_2} u_F(\alpha_F^*(v_2)) \right).$$

Finally, to make the contract Pareto-optimal, weakly Pareto improve each continuation contract to a Pareto-optimal one. □

Proof. For each date 1 state of the world s_1, define $\overline{U}_{1,F}(s_1)$ to be the continuation contract that maximizes F's continuation payoff. Note, in particular, this continuation contract gives nothing to E.

Step 1. Fix a contract and a state $s_1'' = (v_1'', \theta_{1,E}'', \Pi_1'')$ satisfying $U_{1,F}(s_1'') < \overline{U}_{1,F}(s_1'')$. Then for every $v_1' \geq v_1''$, there exists a state $s_1' = (v_1', \theta_{1,E}', \Pi_1')$ satisfying $U_{1,F}(s_1') \leq U_{1,F}(s_1'')$.

Proof of Step 1. Fix a $v_1' \geq v_1''$ and consider the privately observed state $s_1' = (v_1'; \theta_{1,E}', \Pi_1' : (k_1, \delta) \to 2^{\Delta(\{s_2\})})$ with the following properties: $\Pi_1'(k_1) = \Pi_1''([k_1 - (v_1' - v_1'')] \wedge 0)$ and $\Pi_1'(k_1, \delta) = \Pi_1'(k_1, 0)$ for all $\delta \leq v_1' - v_1''$.

Moreover, assume that $\theta_{1,E}'$ small enough such that the continuation contract that maximizes E's continuation payoff subject to delivering continuation payoff $U_{1,F}(s_1'')$ to F satisfies $c_{1,E} + c_{1,F} < v_1''$. That this is possible comes from two observations. First, it is possible to deliver continuation payoff $U_{1,F}(s_1'')$ to F while satisfying $c_{1,F} < v_1''$. To see why, first compute the continuation contract that delivers $\overline{U}_{1,F}(s_1'')$ to F. Obviously, $c_{1,F} \leq v_1''$. Since, by assumption, $U_{1,F}(s_1'') < \overline{U}_{1,F}(s_1'')$, one can just take the continuation contract that delivers $\overline{U}_{1,F}(s_1'')$ to F and decrease $c_{1,F}$ until the continuation payoff decreases to $U_{1,F}(s_1'')$. If after decreasing $c_{1,F}$ all the way to zero the continuation payoff is still larger than $U_{1,F}(s_1'')$, just simply start decreasing date 2 consumption. Second, given that it is possible to deliver continuation payoff $U_{1,F}(s_1'')$ to F while satisfying $c_{1,F} < v_1''$, then by making $\theta_{1,E}$ arbitrarily low, one can ensure that the marginal opportunity cost of consuming at date 1 is arbitrarily high for E. This implies that one can always find a $\theta_{1,E}'$ such that the continuation contract that maximizes E's continuation payoff subject to delivering continuation payoff $U_{1,F}(s_1'')$ to F satisfies $c_{1,E} + c_{1,F} < v_1''$.

Let $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(\cdot), c_{2,F}^*(\cdot)\}$ be the continuation contract that maximizes E's continuation payoff subject to delivering continuation payoff $U_{1,F}(s_1'')$ to F when the state is s_1'.

I claim that $U_{1,F}(s_1') \leq U_{1,F}(s_1'')$.

Suppose not. Suppose s_1' is privately observed by E. If E misreports to v_1'' and $\theta_{1,E}''$ then the reported state is s_1'' and the continuation payoff promised to F is $U_{1,F}''$ which, by assumption, is strictly smaller than what is promised to F if E tells the truth. Now consider the continuation contract $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(\cdot), c_{2,F}^*(\cdot)\}$. If E had told the truth, then F would value this continuation contract at $U_{1,F}(s_1'')$. However, because I assume that $\Pi_1'(k_1, \delta) = \Pi_{1,F}'(k_1, 0)$ for all $\delta \leq v_1' - v_1''$, even if E misreports to v_1'', F's valuation of the continuation contract is unchanged. Thus, by misreporting to s_1'' and then renegotiating the continuation contract to $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(\cdot), c_{2,F}^*(\cdot)\}$, E is made strictly better off and F thinks he is equally well-off. Now just tweak $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(\cdot), c_{2,F}^*(\cdot)\}$ slightly so that F gets slightly more than before, and I have shown that the contract is not renegotiation-proof. Contradiction.

Fix a Pareto-optimal contract and define the following constant:

$$D_1 := \inf_{\{s_1 \ | \ U_{1,F}(s_1) < \overline{U}_{1,F}(s_1)\}} U_{1,F}(s_1)$$

Step 2. E and F weakly prefer the contract described in the theorem with the above D_1 to the Pareto-optimal contract.

Fix a distribution $\pi_1 \in \Pi_1(k_0)$. For every $v_1 < D_1$, move all the weight π_1 puts on v_1-states to a v_1-state where the belief function is trivial. Fix an $\varepsilon > 0$. For every $v \geq D_1$, move all the weight π_1 puts on v-states with $U_{1,F} > D_1 + \varepsilon$ to a v'-state where $U_{1,F} \leq D_1 + \varepsilon$ where $v' \geq v$. Step 1 implies this is possible. Call the modified distribution π_1'. Because the players have lower bound preference-for-robustness, $\pi_1' \in \Pi_1(k_0)$.

Fix any state s with liquidation value $v_1 < D_1$. By assumption, if $U_{1,F}(s) < \overline{U}_{1,F}(s)$ then $U_{1,F} \geq D_1 > v_1$. On the other hand, $\overline{U}_{1,F}(s) \geq v_1$. Thus, $U_{1,F}(s) \geq v_1$. If, furthermore, F's belief function is trivial, then $U_{1,F}(s) = \overline{U}_{1,F}(s) = v_1$. This implies that the value of the portion of the Pareto-optimal contract where $v_1 < D_1$ weakly decreases moving from π_1 to π_1'. It is clear that the value of the portion of the Pareto-optimal contract where $v_1 \geq D_1$ weakly decreases moving from π_1 to π_1'. Thus, the value of the Pareto-optimal contract weakly decreases moving from π_1 to π_1'.

Similarly, the value of the contract described in the theorem weakly decreases moving from π_1 to π_1'. Moreover, the value of the contract described in the theorem plus ε is weakly larger than the value of the Pareto-optimal contract under π_1'. Letting ε tend to zero implies that F weakly prefers the contract described in the theorem to the Pareto-optimal contract.

Next, look at E. For every $v_1 < D_1$, there is a v_1-state s_{v_1} where $U_{1,F}(s_{v_1}) = \overline{U}_{1,F}(s_{v_1})$ and, consequently, E gets nothing. For every $v_1 \geq D_1$, there is a v_1-state s_{v_1} where E gets at most $v_1 - D_1$. Fix a distribution $\pi_1 \in \Pi_1(k_0)$. For every v_1, move all the weight π_1 puts on v_1-states with $U_{1,E} > (v_1 - D_1) \wedge 0$ in the Pareto-optimal contract to s_{v_1}. Call the modified distribution π_1'. $\pi_1' \in \Pi_1(k_0)$. The value of the Pareto-optimal contract weakly decreases moving from π_1 to π_1'. Define a similar modified distribution π_1'' for the contract described in the theorem. The value of the contract described in the theorem weakly decreases moving from π_1 to π_1''. Moreover, the value of the contract described in the theorem under π_1'' is weakly larger than the value of the Pareto-optimal contract under π_1'. Thus, E weakly prefers the contract described in the theorem.

Step 3. The contract described in the theorem is a renegotiation-proof truth-telling contract.

Fix a state with $v_1 < D_1$. Then E gets nothing and the only states that E can (mis)report to are those where the contract maximizes F's continuation payoff subject to E getting nothing. Thus, there is no way he can (mis)report to a state and renegotiate the contract to make himself get more than nothing.

Fix a state with $v_1 \geq D_1$. Then E gets to maximize his payoff subject to delivering D_1 to F. Clearly, E cannot profitably misreport to a state with $\hat{v}_1 < D_1$. If he (mis)reports to any other state, he still has to deliver the same continuation payoff D_1 to F. Moreover, he can only make F weakly more pessimistic about the project by misreporting. Again, there is no way he can (mis)report to a state and renegotiate the contract to make himself strictly better off. □

References

[1] Aghion, P. and P. Bolton (1992) "An Incomplete Contracts Approach to Financial Contracting," *Review of Economic Studies* Vol. 59, pp. 473-494

[2] Carroll, G. (2015) "Robustness and Linear Contracts," *American Economic Review* Vol. 105, pp. 536-563

[3] DeMarzo, P. M. and D. Duffie (1999) "A Liquidity-Based Model of Security Design," *Econometrica* Vol. 67, pp. 65-99

[4] Frankel, A. (2014) "Aligned Delegation," *American Economic Review* Vol. 104, pp. 66-83

[5] Fudenberg, D., Holmstrom, B. and P. Milgrom (1990) "Short-Term Contracts and Long-Term Agency Relationships," *Journal of Economic Theory* Vol. 51, pp. 1-31

[6] Gale, D. and M. Hellwig (1985) "Incentive Compatible Debt Contracts: the One-Period Problem," *Review of Economic Studies* Vol. 52, pp. 647-663

[7] Garrett, D. (2014) "Robustness of Simple Menus of Contracts in Cost-Based Procurement," *Games and Economic Behavior* Vol. 87, pp. 631-641

[8] Gilboa, I. and D. Schmeidler (1989) "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics* Vol. 18, pp. 141-153

[9] Hart, O. and J. Moore (1999) "Foundations of Incomplete Contracts," *Review of Economic Studies* Vol. 66, pp. 115-138

[10] Mookherjee, D. and I. Png (1989) "Optimal Auditing Insurance and Redistribution," *Quarterly Journal of Economics* Vol. 104, pp. 399-415

[11] Mukerji, S. (1998) "Ambiguity Aversion and Incompleteness of Contractual Form," *American Economic Review* Vol. 88, pp. 1207-1231

[12] Roberts, M. (2015) "The Role of Dynamic Renegotiation and Asymmetric Information in Financial Contracting," *Journal of Financial Economics* Vol. 116, pp. 61-81

[13] Roberts, M. and A. Sufi (2009) "Renegotiation of Financial Contracts: Evidence from Private Credit Agreements," *Journal of Financial Economics* Vol. 93, pp. 159-184

[14] Segal, I. (1999) "Complexity and Renegotiation: A Foundation for Incomplete Contracts," *Review of Economic Studies* Vol. 66, pp. 57-82

[15] Townsend, R. (1979) "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory* Vol. 21, pp. 265-293

[16] Yang, M. "Optimality of Debt under Flexible Information Acquisition," working paper