Renegotiation of Dynamically Incomplete Contracts

John Y. Zhu[1]

November 16, 2016

Abstract

I develop a theory of dynamically incomplete financial contracting and explore how to optimally fill in that incompleteness over time. Unlike the related literature, I do not introduce ad-hoc restrictions that make the contracting space exogenously incomplete. By building incompleteness from assumptions on preferences, I go deeper than the incomplete contracts literature and explore *what* can be left incomplete and *when* are control rights needed? This analysis yields the first micro-foundation for debt, defined as a dynamically incomplete contract with a particular preliminary cash-flow right which can then be renegotiated under a default-contingent allocation of control rights.

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

[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, Kenneth Ahern, David Easley, Daniel Ferreira, Alex Frankel, Kris Nimark, Andrew Postlewaite, Robert Ready, Michael Roberts and Maxim Troshkin 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, USC Marshall, University of Rochester Simon.

1 Introduction

Many long-term contracts are dynamically incomplete. For example, a financing contract might officially mature at some date T, and yet the financing relationship will likely extend well beyond that initial maturity date. How exactly should the contract be renegotiated to extend the relationship? These details are often missing. Some contracts like debt allocate state-contingent control rights for filling in these missing details. Other contracts have no instructions at all. In this paper, I develop a theory of dynamically incomplete contracting and explore how that incompleteness should be optimally completed over time.

Unlike most of the incomplete contracting literature, my study of incomplete contracts does not rely on ad-hoc restrictions that make the contracting space exogenously incomplete. By developing incompleteness from the ground up through assumptions about preferences, I am able to study *what* can be left incomplete. From this analysis emerges a deeper understanding of *how* incompleteness should be filled in over time, *when* control rights need to be allocated and *why* they need to be allocated a certain way. These insights about the building blocks of incomplete contracts also yield a richer understanding of specific incomplete contracts like debt and equity. Indeed, the most important contribution I make in this paper is providing a robust micro-foundation for debt as a dynamically incomplete contract featuring default-contingent control rights.

I begin by fixing an ideal dynamic financial contracting setting where the players - an entrepreneur and a financier - can negotiate a complete contract extending to the end of time. No restrictions are placed on the contracting space. Under standard expected utility, it is well known that optimal complete contracts are typically sensitive and complex. Simple, incomplete arrangements like debt are difficult to rationalize. I then ask: What alternative class of preferences would cause optimal complete contracts to always "look like" dynamically incomplete contracts?

Intuitively, if the players have no idea what kind of surplus the relationship can generate in the distant future, they will only agree to a contract that matures before the distant future. Then, over time, if it becomes clear the relationship will continue to generate surplus, the initial maturing contract can then be extended. I formalize this intuition by introducing a class of weakly time-consistent preferences where the players are completely uncertain about the distant future in the Knightian sense. As time passes and distant events move toward the present, that uncertainty is gradually resolved. Players are averse to uncertainty, and, over time, repeatedly make max-min decisions with respect to any residual uncertainty. In my ideal contracting setting, this dynamic max-min preference implies that the payoff generated by an optimal complete contract can be achieved by a dynamically incomplete contract in which the allocation of control rights for filling in missing details is *irrelevant*.

This baseline optimality result makes no predictions about the shape of the initial dynamically incomplete contract or how the missing details should be filled in. It serves as the blank slate on which I further develop my theory of dynamically incomplete contracting. In the rest of the paper, I consider a series of modifications of the baseline setting. I study how these modifications affect both the initial optimal dynamically incomplete contract and the subsequent allocation of control rights. This sheds light on how frictions in the financing relationship force certain features to emerge in the resulting optimal contract such as state-contingent allocation of control rights.

The most important modification I consider is an informational one. Instead of perfect information, I now assume that the entrepreneur privately observes the contract relevant state of the world and can strategically reveal information about it to the financier. I show that an optimal contract responds in two important ways. First, the initial dynamically incomplete contract now gives the financier cash flow rights in a wedge payoff form $D \wedge v$ where D is some constant and v is the value of the underlying asset. Second, the allocation of control rights for filling in missing details is no longer irrelevant and is, instead, done in a state-contingent way with the financier receiving control when $v < D$.

I call this dynamically incomplete contract debt. With respect to the previous literature, my contribution is providing a more *robust* rationale for debt under a more *complete* definition of what is a debt contract.

In terms of robustness, I establish the debt result in a general dynamic complete contracting setting, where players need not be risk-neutral and contracts need not be deterministic. Furthermore, the debt contract is renegotiation-proof when viewed formally as a complete contract. Previous well-known optimality-of-debt results have notably lacked these properties, leading to robustness concerns. For example, Mookherjee and Png (1989) shows that in the costly state verification model, if the financier is allowed to make random verifications, then the optimal contract is no longer debt.

This more robust optimality-of-debt result is also proved under a more complete definition of debt. In my paper debt is a preliminary wedge-shaped cash flow right which can then be renegotiated under a state-contingent allocation of control rights. In contrast, the previous security design literature defines debt largely based on either cash-flow rights or control rights but not really both. For example, there is a complete contracts literature focused on explaining the optimality of wedge-shaped cash flow rights. Seminal papers include Townsend (1979), Gale and Hellwig (1985), and Innes (1990). But those are static models in which there is little scope to talk about control rights ex-post. On the other hand, the primary objective of the incomplete contracts literature is to highlight how variations in the allocation of control rights can have important payoff implications. A classic study of debt using the incomplete contracts approach is Aghion and Bolton (1990). However, there, as in most of the incomplete contracting literature, exogenous restrictions are placed on the contracting space to generate the need for control rights. Moreover, the actions over which control rights are allocated are reduced-form representations of actual contract renegotiations such as restructuring under bankruptcy. In contrast, my approach depends on the dynamic max-min preferences endogenously generating incompleteness and, therefore, a role for control rights. Thus, I can continue to work in a rich, complete contracting setting. In this rich setting, I do not explicitly designate anything as, say, the "bankruptcy restructuring" action. Instead, such a concept emerges organically in my model as an interpretation of the specific way the missing details of a dynamically incomplete contract are optimally filled in over time.

In light of these differences between my debt result and those of the related literature, a key contribution of my paper is providing a robust micro-foundation for debt.

In addition to the modification of the baseline setting leading to debt, I consider another modification concerning how uncertainty about distant events is resolved as those events move toward the present. I specialize to a setting where the interim uncertainty takes a form where the set of beliefs the players are willing to entertain consists of "towers" of beliefs ordered by the mean-preserving spread relation. A canonical example is when the players have beliefs about the asset's expected value but do not know the precise distribution. In this case I show that optimal dynamically incomplete contracts involve players beginning with linear shares of the asset and renegotiating to new linear shares over time as conditions change. This motivates equity-based dynamic financing arrangements.

Again, an important quality of the result is robustness - players need not be risk-neutral and contracts need not be deterministic. Moreover, I establish the result in a dynamic setting. Most previous linearity results are about static contracts and the few that do deal with dynamic contracts require strong assumptions on fundamentals and involve each party receiving a permanently fixed linear share. In contrast, my linearity result yields a more flexible notion of contracting with equity in a general dynamic contracting setting. My dynamic linearity result also links this paper to a related static contracting literature seeking to micro-found the null contract in the classic hold-up model. See, for example, Che and Hausch (1999), Hart and Moore (1999), and Segal (1999). The closest paper is Mukerji (1998) which points out that the null contract, by giving each party a fixed fraction of the future surplus, may be attractive to max-min decision makers. The idea anticipates the one behind my linearity result and is related to the insight of Carroll (2015).

2 A Toy Model

In the introduction I explained that I will first establish in a baseline setting a "blank slate" result about the optimality of dynamically incomplete contracts where the incompleteness emerges endogenously from assumptions about preferences. Then I will modify the baseline setting in various ways and see how optimal contracts respond, yielding, among other things, a micro-foundation for debt. In this section, I partially carry out that program in a toy model analysis. I introduce a toy version of the baseline setting and establish the optimality of dynamically incomplete contracts. Then, I perform the second modification described in the introduction on the toy baseline setting and show that these optimal contracts become linear.

There are two players, E and F, who share an asset spanning dates 1 and 2. At date 1, the asset produces some random amount of capital $v_1 \sim U[0, x_1]$. A portion of this capital can be withdrawn for consumption, $c_{1,E}$ for E and $c_{1,F}$ for F. The asset then turns the remaining capital $k_1 := v_1 - (c_{1,E} + c_{1,F})$ into some random amount of capital $v_2 \sim U[0, x_2 k_1]$ at date 2 which is then consumed by the players, $c_{2,E} + c_{2,F} = v_2$. x_1 is known at the time of contracting which occurs before date 1 while x_2 is realized at date 1.

A complete contract specifies a history-dependent consumption plan

$$\{c_{1,E}(v_1, x_2), c_{1,F}(v_1, x_2), c_{2,E}(v_1, x_2, v_2), c_{2,F}(v_1, x_2, v_2)\}.$$

A belief is completely summarized by a joint distribution π_1 of (v_1, x_2) with the property that $\pi_1|_{v_1} \sim U[0, x_1]$. Here, $|_{v_1}$ means the marginal distribution over v_1. The optimal contracting problem is to characterize the set of all renegotiation-proof contracts.

Definition. *Given a contract, a date 1 continuation contract is the contract with* (v_1, x_2) *treated as fixed. A contract is renegotiation-proof at date 1 if every date 1 continuation contract is Pareto-optimal. A contract is renegotiation-proof if it is Pareto-optimal among all contracts that are renegotiation-proof at date 1.*

The standard expected utility setting is one where E and F share a single belief π_1, so that E (and similarly F) evaluates contracts as follows,

$$\mathbf{E}_{(v_1,x_2)\sim\pi_1}\left[u_E(c_{1,E}(v_1,x_2)) + \mathbf{E}_{v_2\sim U[0,x_2k_1]}u_E(c_{2,E}(v_1,x_2,v_2))\right].$$

Here, u_E is the weakly concave, strictly increasing utility function of E. It is well known that optimal dynamic contracts under expected utility are typically complex and not robust. Moreover, they do not look dynamically incomplete. In contrast, my goal is to make robust predictions about dynamically incomplete contracts.

To achieve my goal, I now introduce a toy version of the baseline setting which features the dynamic max-min preferences that I discussed in the introduction and will formally define in the next section. At the time of contracting, the players have no idea if the relationship can generate any surplus in the distant future which is represented by date 2. That is, the players not only entertain belief π_1 but also any other belief π_1' satisfying $\pi_1'|_{v_1} = \pi_1|_{v_1} = U[0,x_1]$. Suppose E and F are averse to their inability to pin down a precise belief. Then E (and similarly F) evaluates contracts as follows,

$$\min_{\left\{\pi_1' \;\middle|\; \pi_1'|_{v_1}=U[0,x_1]\right\}} \mathbf{E}_{(v_1,x_2)\sim\pi_1'}\left[u_E(c_{1,E}(v_1,x_2)) + \mathbf{E}_{v_2\sim U[0,x_2k_1]}u_E(c_{2,E}(v_1,x_2,v_2))\right].$$

Theorem 1. *Every Pareto-optimal payoff can alternatively be achieved by a dynamically incomplete contract with unspecified renegotiation.*

Here, the Pareto-frontier is the payoff frontier generated by renegotiation-proof complete contracts. I will define what exactly is a "dynamically incomplete contract with unspecified renegotiation" in the course of proving the result.

Theorem 1 is the "blank slate" I referred to in the introduction. As will become clear from the proof, the result makes no predictions about the shape of the dynamically incomplete contract or how it should be renegotiated to fill in missing details. The rest of the paper is then focused on modifying the baseline setting in various ways so that I can say something more specific about the structure of optimal contracts.

Proof. Fix a point on the Pareto-frontier. There are many renegotiation-proof complete contracts that achieve this payoff point. The entire collection is a union of sets, each of which can be characterized as follows: There exists a split $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ of v_1 depending only on v_1 such that a contract belongs in the set if and only if it is renegotiation-proof at date 1 and every date 1 continuation contract Pareto-dominates the split - that is,

$$\begin{aligned}\big(u_E(c_{1,E}(v_1,x_2)) + \mathbf{E}_{v_2\sim U[0,x_2k_1]}u_E(c_{2,E}(v_1,x_2,v_2)),\\ u_F(c_{1,F}(v_1,x_2)) + \mathbf{E}_{v_2\sim U[0,x_2k_1]}u_F(c_{2,F}(v_1,x_2,v_2))\big)\\ \geq (u_E(\alpha_E^*(v_1)), u_F(\alpha_F^*(v_1)))\end{aligned}$$

for every (v_1, x_2). This characterization of the set admits the following interpretation:

Consider a tentative split $(\alpha_E(v_1), \alpha_F(v_1))$ of v_1 representing outside options for E and F at date 1. This split does not necessarily equal actual date 1 consumption. Instead, it serves as the individual rationality constraints for E and F at date 1 as they renegotiate a Pareto-optimal continuation contract. Let us refer to this split as a *dynamically incomplete contract* since there are no instructions for how the split should be renegotiated at date 1.

Under expected utility, players cannot rank dynamically incomplete contracts and therefore, it is not well-defined to talk about optimal dynamically incomplete contracting. The characterization of the set of renegotiation-proof complete contracts that achieve a common payoff implies that under the current max-min setup, optimal dynamically incomplete contracting is well-defined. Given an arbitrary split $(\alpha_E(v_1), \alpha_F(v_1))$, so long as it is always Pareto-improved to *some* Pareto-optimal date 1 continuation contract, the players can assign a value to the split while still being agnostic about *which* Pareto-optimal date 1 continuation contract will be chosen.

Thus, to achieve a point on the complete contracts Pareto-frontier, it suffices for E and F to initially negotiate only an optimal dynamically incomplete contract $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ that leaves details concerning any future renegotiation completely unspecified. Then, once date 1 arrives, E and F can renegotiate freely to some Pareto-optimal date 1 continuation contract that Pareto-dominates the split. □

These optimal dynamically incomplete contracts lack any ex-ante restrictions on how future renegotiations should be settled. This open-ended property is sometimes referred to as "ex-post renegotiation." Despite the intended meaning of ex-post renegotiation, much of the previous literature models it as a specific renegotiation protocol that gives each party a constant share of any surplus generated from renegotiation. Players are assumed to know, ex-ante, what their constant shares will be. My concept of "unspecified renegotiation" better captures the we'll-cross-that-bridge-when-we-get-there spirit of ex-post renegotiation.

The fact that contracts can allow for unspecified renegotiation and still be optimal depends, in part, on the perfect information assumption. If, instead, E were to privately observe (v_1, x_2) then, depending on the shape of the initial dynamically incomplete contract, there could be scope for manipulation and the financing relationship could be compromised. In a few special cases this is not a problem. For example, if it is known ex-ante that v_1 will never be lower than some constant $\underline{v}$, then a dynamically incomplete "safe debt" contract like $(\alpha_E^*(v_1) = v_1 - \underline{v}, \alpha_F^*(v_1) = \underline{v})$ with unspecified renegotiation at date 1 is still robust to manipulation.

What can the parties do when safe debt contracts do not exist? In the formal analysis of the next section, I tackle this question and show that optimal complete contracts look like dynamically incomplete contracts where the renegotiation at date 1 is partially specified through the allocation of state-contingent control rights.

Before moving on to that formal analysis, I perform a modification on the toy baseline setting - the second one described in the introduction - and show that it results in optimal dynamically incomplete contracts that are linear.

Linear Dynamically Incomplete Contracts

Recently, there has been significant progress made in providing a robust rationale for linear contracts in the static contracting setting. Here, robustness means that the optimality of linear contracts does not depend on very specific assumptions about the distributional nature of the underlying uncertainty. Building on Mukerji (1998) and Carroll (2015), I strengthen the robustness of the linearity result in two directions. First, I relax the requirement that contracts must be deterministic and players must be risk-neutral. Second, I move from the static setting into the dynamic setting.

In the toy baseline setting considered above, uncertainty about the distant future resolves itself in a rather stark way: At date 0, players have no idea how v_2 will be distributed. One date later, they have a precise belief that $v_2 \sim U[0, x_2 k_1]$. I now modify the setting so that at each date t, the players continue to have residual uncertainty about v_{t+1}.

Specifically, at date 1, players not only entertain the belief $v_2 \sim U[0, x_2 k_1]$, but, more generally, any belief π_2 such that $\mathbf{E}_{\pi_2} v_2 = \frac{x_2 k_1}{2}$. This captures a situation where players know the expected value but nothing beyond that. Moving back to the moment of contracting before date 1, the players similarly not only entertain the belief π_1, but, more generally, any belief π_1' over (v_1, x_2) with the property that $\mathbf{E}_{\pi_1'} v_1 = \frac{x_1}{2}$. I show later that the linearity result holds under more general assumptions about what beliefs E and F can entertain. For now, I assume this particular structure for the sake of concreteness. Given this modification, E (and similarly F) now evaluates contracts as follows

$$\min_{\left\{\pi_1' \;\middle|\; \mathbf{E}_{\pi_1'} v_1 = \frac{x_1}{2}\right\}} \mathbf{E}_{(v_1, x_2) \sim \pi_1'} \left[u_E(c_{1,E}(v_1, x_2)) + \min_{\left\{\pi_2 \;\middle|\; \mathbf{E}_{\pi_2} v_2 = \frac{x_2 k_1}{2}\right\}} \mathbf{E}_{v_2 \sim \pi_2} u_E(c_{2,E}(v_1, x_2, v_2)) \right].$$

Corollary 1. *Every Pareto-optimal payoff can alternatively be achieved by a linear dynamically incomplete contract with unspecified linear renegotiation.*

What this result says is that every point on the Pareto-frontier can be achieved by a dynamically incomplete contract with unspecified renegotiation where $(\alpha_E^*(v_1), \alpha_F^*(v_1)) = (\alpha_E^* \cdot v_1, \alpha_F^* \cdot v_1)$ for some constants $\alpha_E^*, \alpha_F^* = 1 - \alpha_E^*$. Moreover, every Pareto-optimal date 1 continuation contract is payoff equivalent to one that specifies a linear split of v_2. Thus, at date 1, when the players renegotiate $(\alpha_E^* \cdot v_1, \alpha_F^* \cdot v_1)$ to some individually rational Pareto-optimal continuation contract, it is without loss of generality to assume that they renegotiate to one that splits v_2 linearly. I call this linear renegotiation.

An immediate implication of Corollary 1 is that, instead of renegotiating over the entire space of dynamically incomplete contracts, the players can simply create equity claims on the asset and trade them. The players begin with equity stakes (α_E^*, α_F^*). Then at date

1, these claims on the asset establish the players' bargaining positions as they negotiate dividend payouts (i.e. $c_{1,E}$ and $c_{1,F}$) and their future equity stakes in the asset.

Proof. The proof is implied by analyzing a related static risk-sharing problem. Suppose E and F share an asset that produces a one time random payoff v. Ex-ante, for every Pareto-optimal payoff, the players seek to find a sharing rule $(\alpha_E(v), \alpha_F(v))$ that will achieve it. Under expected utility, a well-known necessary condition for Pareto-optimality is *Borch's Rule*, which states that if $(\alpha_E(v), \alpha_F(v))$ is Pareto-optimal then the ratio $u'_E(\alpha_E(v))/u'_F(\alpha_F(v))$ of marginal utilities must be constant across all v. Borch's Rule implies that, outside of a few special cases, linear sharing rules are suboptimal.

Now suppose E and F are uncertain about v in the way described above and are averse to uncertainty. That is, there is some constant $\overline{v}$ such that they entertain any belief π satisfying $\mathbf{E}_\pi v = \overline{v}$ and evaluate shares as follows,

$$\left(\min_{\{\pi \ | \ \mathbf{E}_\pi v = \overline{v}\}} \mathbf{E}_{v\sim\pi} u_E(\alpha_E(v)), \ \min_{\{\pi \ | \ \mathbf{E}_\pi v = \overline{v}\}} \mathbf{E}_{v\sim\pi} u_F(\alpha_F(v)) \right).$$

Consider E's share $\alpha_E(v)$. Suppose it is strictly concave over, say, $[0,1]$. Fix a distribution π that E entertains and suppose it puts some weight on the open interval $(0,1)$. Then construct an alternative distribution π' which is identical to π except that it shifts the weight on $(0,1)$ to the endpoints in a mean-preserving way. Since E entertains π, he entertains π' as well. Moreover $\mathbf{E}_{v\sim\pi} u_E(\alpha_E(v)) \geq \mathbf{E}_{v\sim\pi'} u_E(\alpha_E(v))$. Given E's max-min preferences, this observation implies that it is without loss of generality for E to entertain only those distributions that do not put any weight on the strictly concave regions of $\alpha_E(v)$. But then it is a weak Pareto-improvement to linearize all the strictly concave regions of $\alpha_E(v)$ and give the surplus to F. This implies that it is without loss of generality to assume that $\alpha_E(v)$ is weakly convex. By the same reasoning, it is without loss of generality to assume $\alpha_F(v)$ is weakly convex. The only sharing rules that are weakly convex for both players are precisely the linear sharing rules. $\square$

The proof implies that the linearity result holds under the more general assumption that each player entertains a possibly distinct set of beliefs consisting of towers of beliefs ordered by the mean-preserving spread relation. Such a tower is defined to consist of a most precise belief π and all beliefs that are mean-preserving spreads of it. One caveat is that players need to become risk-neutral when v_t is sufficiently large. Otherwise, the Pareto-frontier collapses to a point and the linearity result has no bite. There are also various other ways to get the optimality or near-optimality of linear contracts without assuming risk-neutrality over high values of v_t by slightly modifying the structure of the set of beliefs E and F entertain. The simplest way is to assume that v_t is bounded above and that all beliefs must respect that bound.

The proof also implies that the linearity result is essentially a result about static contracts. Corollary 1 simply shows how this static result can be imported into a dynamic setting. In contrast, the debt result I prove in the next section is truly a result about dynamic contracts. The wedge-shaped cash flow rights and the default-contingent control rights that comprise debt are jointly optimal in the dynamic asymmetric information setting. There is no static version of that setting in which I can exclusively derive the wedge-shaped cash flow rights as the optimal static contract.

3 The Formal Model

In this section, I first introduce the formal baseline setting and re-establish the "blank slate" optimality result, Theorem 1. Then I modify the setting by introducing asymmetric information and prove the optimality of debt. The need to move beyond the toy model and to introduce the formal model arises due to the debt analysis. As will become clear, the optimality of debt hinges on the financier worrying about many different things that can "go wrong." The toy model, with its simplifying assumptions, simply removes too many degrees of freedom for debt to emerge as the optimal dynamically incomplete contract.

Once again, E and F share an asset spanning dates 1 and 2. In addition, there is an ex-ante date 0. At date 1, the asset produces some random amount of capital $v_1 \geq 0$. The players consume $c_{1,E} + c_{1,F} \leq v_1$. The remaining capital, call it k_1, then randomly generates an amount of date 2 capital v_2 which is completely consumed, $c_{2,E} + c_{2,F} = v_2$.

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,

$$s_1 := (v_1, \Pi_2 : k_1 \to 2^{\Delta(\{s_2\})}).$$

A date 1 state of the world consists of a realized date 1 capital plus a *belief function* Π_2, which specifies for each possible date 1 remaining capital $k_1 \in [0, v_1]$ a set of beliefs about the date 2 state of the world. I impose some mild restrictions on the shape of Π_2. I assume Π_2 is weakly increasing - that is, 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. Fix a subset $\mathcal{S}_1 \subset \{s_1\}$. Given $\mathcal{S}_1$, an ex-ante date 0 state of the world $s_0 := (\Pi_1 \in 2^{\Delta(\mathcal{S}_1)})$ is a set of beliefs about the date 1 state of the world subject to the restriction that all beliefs have supports that are subsets of $\mathcal{S}_1$.

Notice I assume that players share the same belief function at all times. I develop the model this way purely for notational simplicity. The model can be reformulated so that states of the world specify potentially distinct belief functions for the players. All results would continue to hold.

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

Contracts and Dynamic Max-Min Preferences. Fix a setting $(s_0, \mathcal{S}_1)$. A complete contract is a history-dependent consumption plan for E and F,

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

A contract implies a state-contingent remaining capital amount $k_1(s_1)$.

As I have defined them, contracts are not random - consumption can only depend on the history of the state of the world, and not some exogenous randomizing devices. However, all of my results continue to go through even allowing for random contracts.

Notice, the definition of contracts implies that belief functions Π_t are contractible. To the extent that these Π_t are formulated based on publicly observable data about the asset (e.g. credit rating) or the market (e.g. VIX), this can be justified. Realistically though, such belief functions are at least partially based on unobservable or unverifiable information and one could argue that therefore they should not be modeled as contractible. However, since I am analyzing a perfect information benchmark right now and since my goal is to provide a rationale for why contracts are dynamically incomplete, I want to place as few exogenous restrictions on the contracting space as possible. By letting literally everything be contractible, I create the greatest distance between myself and my goal, allowing me to better develop a non-ad-hoc theory of optimal dynamically incomplete contracting. Of course, once I introduce asymmetric information, I will no longer allow the privately observed state of the world to be directly contractible. However, I will still allow E to make contractible reports about the state.

Given a contract (c_E, c_F), E's continuation payoff process U_E characterizing his preferences is defined as follows,

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

U_F is defined similarly for F.

Lemma 1. *The dynamic max-min preference is weakly time-consistent.*

Proof. See Appendix. □

By selecting the appropriate setting $(s_0, \mathcal{S}_1)$, this dynamic max-min preference can reduce to standard expected utility (when $(s_0, \mathcal{S}_1)$ features only singleton-valued belief functions) or the preferences used to develop my theory of dynamically incomplete contracts.

Recall, in the introduction I outlined the following program: 1. Introduce a baseline setting in which optimal complete contracts can be implemented by dynamically incomplete contracts in which the allocation of control rights is irrelevant. 2. Perform a series of modifications on the baseline setting and see how optimal contracts respond.

In the previous toy model section, I partially carried out this program. I introduced a toy version of the baseline setting, established the blank slate optimality result, Theorem 1), and performed one modification of the baseline setting that led to the optimality of linear dynamically incomplete contracts. All of this can be re-done in the formal model.

The formal baseline setting is the following $(s_0, \mathcal{S}_1)$,

Assumption. $\mathcal{S}_1 = \{s_1\}$. *s_0 is a date 0 state of the world satisfying the following property: There is a set of distributions V_1 of v_1 such that $\pi_1 \in \Pi_1$ if and only if $\pi_1|_{v_1} \in V_1$.*

The statement and interpretation of the baseline setting's optimality result are unchanged.

Theorem 1. *Every Pareto-optimal payoff can alternatively be achieved by a dynamically incomplete contract with unspecified renegotiation.*

Moreover, the formal baseline setting can be modified in a way similar to what was done in the toy section and Corollary 1 can be re-established. The proofs are essentially the same so I will not dwell on these results anymore.

Instead, I now focus on the main modification of the baseline setting described in the introduction - replacing perfect information with asymmetric information about the state of the world. I will then establish the main result concerning the optimality of debt.

3.1 Asymmetric Information and Debt

In a series of papers, Mukerji (1998), Che and Hausch (1999), Hart and Moore (1999), and Segal (1999) provide a micro-foundation for the null incomplete contract. This involves proving that the optimal complete contract in a reasonably general contracting model is equivalent to the null incomplete contract. In this section I do the same for the debt incomplete contract.

I now assume that the date 1 state of the world is privately observed by E and contracts can depend on s_1 only through a report $\hat{s}_1$ made by E. For simplicity, I continue to assume that the date 0 and date 2 states of the world are observable.

The publicly observed date 2 state of the world still consists of a realized v_2. The privately observed date 1 state of the world s_1 is the following object,

$$s_1 := (v_1, \theta_{1,E}, \Pi_2 : (\hat{k}_1, \delta) \to 2^{\Delta(\{s_2\})}).$$

Lastly, given $\mathcal{S}_1 \subset \{s_1\}$, a publicly observed date 0 state of the world is still some $s_0 := (\Pi_1 \in 2^{\Delta(\mathcal{S}_1)})$.

Notice, the privately observed s_1 is a different object compared to the s_1 under perfect information. There is a new parameter $\theta_{1,E}$ and the "belief function" Π_2 is now a function of a pair of as yet undefined parameters $\hat{k}_1$ and δ. For now, I set Π_2 aside.

The parameter $\theta_{1,E}$ is a transitory taste shock for E. Given a realized $\theta_{1,E} \in \mathbb{R}^+$, E's utility of consumption for date 1 is $\theta_{1,E} u_E(\cdot)$. While this taste shock did not appear in the baseline setting, the exclusion was purely for simplicity - Theorem 1 would have continued to hold with taste shocks. So why do I introduce taste shocks now?

A taste shock at date 1 affects E's intertemporal rate of substitution. Loosely speaking, changing this intertemporal rate of substitution changes E's opinion about the best way to deploy the asset's capital at date 1. For example, if $\theta_{1,E}$ is very high, then E's opinion is that saving capital to invest in the future is inferior to consuming the capital today.

It turns out, F's belief about what E's opinion might be plays an important role in the optimality of debt. Specifically, a key ingredient in the proof is the combination of, ex-ante, F being completely uncertain about E's future opinion, and, ex-post, F not observing that opinion. Introducing a hidden taste shock to the setting is simply a convenient way of generating this key ingredient. Of course, taste shocks are not the only way. I could have also generated the key ingredient by directly assuming that E and F can have differing belief functions. I choose the taste shock route mostly for simplicity.

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

Let us now return our attention to the object $\Pi_2(\hat{k}_1, \delta)$.

In short, $\Pi_2(\hat{k}_1, \delta)$ tells E not only what his belief function is but also what F's belief function will be as a function of how he reports the state.

More specifically, let us assume that F expects E to report truthfully - this will be justified by my restricting attention to truth-telling contracts later on. Then, on the equilibrium path where E does indeed tell the truth, F should have the same belief function as E. This is in keeping with the perfect information baseline setting's simplifying assumption that both parties share the same belief function. Let us call this common belief function $\Pi_{2,\delta=0} : k_1 \to 2^{\Delta(\{s_2\})}$.

In the current asymmetric information setting, a date 1 state of the world should not only specify $\Pi_{2,\delta=0} : k_1 \to 2^{\Delta(\{s_2\})}$ but should also model how F's belief function might change given off-equilibrium path misreports by E. For simplicity, assume that F's belief function is affected only by misreports of v_1, not $\theta_{1,E}$. Then, for each δ, the date 1 state of the world should specify a $\Pi_{2,\delta} : \hat{k}_1 \to 2^{\Delta(\{s_2\})}$ that would be F's off-equilibrium belief function if E deviated from the equilibrium path and reported $\hat{v}_1$. This collection of belief functions is captured by the two-dimensional function $\Pi_2 : (\hat{k}_1, \delta) \to 2^{\Delta(\{s_2\})}$ that I include as part of the date 1 state of the world.

The way I have modeled asymmetric information endows E with great powers of perception. F is facing a E who can finely tailor his reporting strategy to take advantage of any perfectly perceived flaw in F's ability to forecast the future productivity of the asset. As a result, F will demand a contract that is robust enough to defend against such a powerful E. In practice, the disparity between E and F may not be so dramatic. Thus, one can view the optimal contract I derive in this section as a robust lower bound contract that F can always fall back on as he contemplates potentially better alternatives.

I impose two mild restrictions on the shape of $\Pi_2(\hat{k}_1, \delta)$. For every δ, Π_2 is weakly increasing in $\hat{k}_1$ with $\Pi_2(0, \delta) = 0$ just like in the perfect information baseline setting. Also, I assume that Π_2 is weakly decreasing in δ holding k_1 fixed. This means the more E underreports v_1, the weakly more pessimistic is F about the asset's future productivity.

Given reports $(\hat{v}_1, \hat{\theta}_{1,E})$, the induced public reported date 1 state of the world is

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

A contract specifies a history-dependent consumption plan

$$(c_E, c_F) := (c_{1,E}(\hat{s}_1), c_{1,F}(\hat{s}_1), c_{2,E}(\hat{s}_1, s_2), c_{2,F}(\hat{s}_1, s_2))$$

along with a report strategy $(\hat{v}_1(s_1), \hat{\theta}_{1,E}(s_1))$ for E. A contract is truth-telling if $\hat{v}_1 \equiv v_1$

and $\hat{\theta}_{1,E} \equiv \theta_{1,E}$. As an abuse of notation, I let $s_1 = (v_1, \theta_{1,E}, \Pi_2 : k_1 \to 2^{\Delta(\{s_2\})})$ denote the reported state when the privately observed state is $s_1 = (v_1, \theta_{1,E}, \Pi_2 : (\hat{k}_1, \delta) \to 2^{\Delta(\{s_2\})})$ and E tells the truth.

I focus on renegotiation-proof truth-telling contracts. Under full commitment, restricting attention to truth-telling contracts is without loss of generality - this is the revelation principle. However, when commitment is relaxed, which is the case when doing renegotiation-proof contracting, it is with loss of generality to restrict attention to truth-telling contracts. In fact, renegotiation-proof truth-telling contracts are not even guaranteed to exist in general.

That being said, given the intrinsic appeal of the full information revelation implied by truth-telling, if a contracting setting did happen to have renegotiation-proof truth-telling contracts, then it would be worth figuring out what they were, independent of any analysis of a more general optimal contracting problem that allows pooling of types. As it turns out, in my setting, renegotiation-proof truth-telling contracts do exist and as I will prove shortly, they are precisely the debt 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 feasible (possibly off-equilibrium) reported state $\hat{s}_1$. The continuation contract is some

$$(c_{1,E}(\hat{s}_1), c_{1,F}(\hat{s}_1), c_{2,E}(\hat{s}_1, s_2), c_{2,F}(\hat{s}_1, s_2)).$$

where $\hat{s}_1$ is now treated as fixed. A *renegotiation* of this continuation contract is some alternate consumption plan

$$(c'_{1,E}, c'_{1,F}, c'_{2,E}(s_2), c'_{2,F}(s_2))$$

satisfying the same date 1 budget constraint $c'_{1,E} + c'_{1,F} \leq \hat{v}_1$. The true remaining capital is now $k'_1 = v_1 - (c'_{1,E} + c'_{1,F})$ while the reported remaining capital is $\hat{k}'_1 = \hat{v}_1 - (c'_{1,E} + c'_{1,F})$.

Under the renegotiation, E's continuation payoff is

$$u_E(c'_{1,E}) + \min_{\pi_2 \in \Pi_2(k'_1)} \mathbf{E}_{s_2 \sim \pi_2} u_E(c'_{2,E}(s_2)), \tag{1}$$

and F thinks his continuation payoff is

$$u_F(c'_{1,F}) + \min_{\pi_2 \in \Pi_{2,\delta}(\hat{k}'_1)} \mathbf{E}_{s_2 \sim \pi_2} u_F(c'_{2,F}(s_2)). \tag{2}$$

E is strictly better off under the renegotiation if

$$(1) > u_E(c_{1,E}(s_1)) + \min_{\pi_2 \in \Pi_2(k_1)} \mathbf{E}_{s_2 \sim \pi_2} u_E(c_{2,E}(s_1, s_2)),$$

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

$$(2) > u_F(c_{1,F}(\hat{s}_1)) + \min_{\pi_2 \in \Pi_{2,\delta}(\hat{k}_1)} \mathbf{E}_{s_2 \sim \pi_2} u_F(c_{2,F}(\hat{s}_1, s_2)).$$

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 he is strictly better off.*

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

Definition. *A debt contract is a dynamically incomplete contract* $(\alpha_E^*(v_1), \alpha_F^*(v_1))$ *of the form* $(v_1 - D_1 \wedge v_1, D_1 \wedge v_1)$ *for some constant* D_1 *with partially specified renegotiation: F (E) is allocated the control right to choose the individually rational continuation contract at date 1 if* $v_1 < D_1$ $(v_1 \geq D_1)$*.*

In the security design literature, it is common to impose some monotonicity constraint on the contracting space in order to prove the optimality of debt. I could do that here as well, but I do not need to if, instead, I modify the baseline setting $(s_0, \mathcal{S}_1)$ slightly: Replace the condition $\pi_1|_{v_1} \in V_1$ with $\pi_1|_{v_1} \geq_d \mu_1$ for some $\mu_1 \in V_1$. I now assume this modified baseline setting.

Theorem 2. *Every Pareto-optimal payoff can alternatively be achieved by debt.*

The formal meaning of the theorem is the following: Assuming players act optimally when allocated control rights, a debt contract combined with the truth-telling report strategy is the same object as a particular complete contract. This complete contract is renegotiation-proof. Moreover, every point on the Pareto-frontier can be achieved by such a complete contract.

To get a feel for how the proof of Theorem 2 works, let us take the debt contract, keep the initial split the same but change the partial specification of renegotiation so that E is allocated control all the time. I will now informally argue that this is a poor contract for F.

First, it is important to remember that, in general, the individually rational date 1 continuation contract E wants to select 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 occurred 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 about future productivity.

Then, I claim that E is better off 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 weakly 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 weakly less, he may still have to give up strictly more to deliver that smaller continuation payoff due to the increased pessimism.

But notice, neither of these two potential downsides are a concern for us. I have assumed a state where E does not want to consume much at date 1 anyways and I 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. Here is where F's ex-ante uncertainty about the future state of the world matters. Even though the "worst case" scenario of underreporting does not always occur, it matters at the margin for someone who is uncertainty-averse.

My informal analysis of the altered debt contract suggests the following lesson,

Remark. *Varying F's payoff with asset 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 2 is basically a more fleshed out version of the arguments made here.

Just how bad can things get if, in the default region, control is taken away from F?

Corollary 2. *Fix an arbitrary debt contract. Change the partially specified renegotiation so that when in the default region, F is able to choose any individually rational contract that gives himself up to some constant fraction of the surplus that is strictly smaller than one. Then the modified debt contract sometimes does not induce truth-telling. Moreover, the ex-ante payoff of the contract to F is zero.*

Proof. This is implied by Step 2 of the proof of Theorem 2. □

This micro-foundation for debt also provides some formal justifications for other arrangements that have been studied in the incomplete contracts literature. For example, Chung (1991) studies a variable quantity procurement model and focuses attention on a class of simple incomplete contracts. A contract in that class consists of an initial quantity-price pair followed by a revision scheme granting either the buyer or the seller the right to make a take-it-or-leave-it offer ex-post consisting of a new quantity-price pair. That incomplete contracting convention closely matches the one that endogenously emerges in my paper. Recall, my initial dynamically incomplete contract is the outside option in the subsequent renegotiation. It serves the same function as the initial quantity-price pair in Chung's model. Also, my partially specified renegotiation allocates control rights for choosing a new individually rational continuation contract to replace the initial one. This mirrors Chung's revision scheme which allocates bargaining power in choosing a new quantity-price pair to replace the initial one.

One concern with regards to Chung's incomplete contracts is that there is a tension between imposing the renegotiation-proof condition on contracts and allowing contracts to allocate bargaining power. That is, does it make sense to assume that the players can commit to play the take-it-or-leave-it bargaining game that the contract requires? In my debt contract, this is technically not an issue because E and F are technically not playing a bargaining game. Instead, one player is simply choosing the individually rational continuation contract. The distinction emerges because the contractible reported state in my model contains all information needed to determine if a chosen continuation contract is individually rational. Thus, for example, in the non-default region, if E chooses an individually rational contract and F falsely claims that it is below his outside option and rejects it, E can prove that F has lied. This, of course, rests on the assumption that the induced belief function of F is part of the contractible reported state. One might rightly wonder then: How robust is my control rights arrangement to changes in the verifiability of F's belief function?

It turns out, I can successfully implement my control rights arrangement even without F's belief function being verifiable so long as there is a competitive financing market in which there is another $\hat{F}$ who is identical to F and who is willing to step in and provide funding. In this case, simply modify the debt contract so that E can unilaterally ask $\hat{F}$ to buy out F's position. Here, unilateral means subject to $\hat{F}$'s approval but not F's. A buyout request from E consists of having $\hat{F}$ pay $\alpha_F^*(v_1)$ to F in exchange for some continuation contract $(c_{1,E}, c_{1,\hat{F}}, c_{2,E}(s_2), c_{2,\hat{F}}(s_2))$.

The implementation of the optimal contract by debt differs from the previous implementation results in one important way: Renegotiation can no longer be completely unspecified. Contrasting Theorem 2 with Theorem 1 highlights how the state-contingent allocation of control rights emerges as an optimal response to information asymmetry. Without this information asymmetry, then once again the allocation of control rights would be irrelevant. The initial wedge-shaped cash flow rights could be renegotiated in a completely unrestricted way and the ex-ante payoffs of the dynamically incomplete contract to both parties would be unaffected.

Of course, without information asymmetry, there would also be no compelling reason for the initial dynamically incomplete contract to take a wedge-shaped form. Thus, in my model, information asymmetry simultaneously necessitates wedged-shaped cash flow rights and default-contingent control rights. This rationale for why the two salient features of debt are jointly optimal, which applies in a general dynamic complete contracts model, is a key contribution of my paper.

4 Future Directions

My work has built on a number of recent papers looking at static contracting under max-min preferences. See, for example, Antic (2014), Frankel (2014), Garrett (2014), and Carroll (2015). A related approach looks at contracting under standard expected utility where the incentive problems have a similarly flexible nature but costs are parameterized using an entropy function. See Hebert (2015) and Yang (2015). In this paper I extended the static max-min approach to a dynamic setting. I introduced a baseline optimal contracting setting under dynamic max-min preferences and then studied two modifications of the setting and their effects on the structure of optimal dynamically incomplete contracts. There are obviously many more directions to investigate, and some of them may lead to new insights.

Take, for example, the debt contract defined in the previous section that allocates F (E) control in default (non-default) states. Corollary 2 implies that changing the default region control rights even slightly can have serious negative effects on F's ex-ante payoff. What about changing non-default region control rights? Such changes can certainly induce E to misreport but E will never misreport to a default state. This means F's ex-post payoff can only weakly increase and, in fact, ex-ante payoffs are unaffected. Perhaps then there is a way of relaxing the truth-telling constraint so that debt remains optimal, but renegotiation in the non-default region can be left completely unspecified. This would then introduce room to think about how additional frictions in the model might map to different partial specifications for renegotiation in the non-default region.

Another direction to investigate involves stretching out the horizons of the players. Right now, players only have some idea about the asset's value tomorrow. Letting players not be completely uncertain about events two or more dates into the future could introduce further subtleties that are hard to capture in the current setup. With longer horizons, there would also be many more ways to model how uncertainty about distant events gets gradually resolved over time.

A thorough analysis of optimal contracting along these lines is beyond the scope of my paper. However, I consider one simple example where players' horizons are longer. Working through this example will serve to both explicitly demonstrate one way my class of dynamic max-min preferences can model longer horizons and provide a glimpse of the new results that may come from such an analysis.

4.1 Interim Renegotiation

So far I have shown how optimal complete contracts can be implemented as dynamically incomplete contracts followed by renegotiation at maturity. However, empirical work by Roberts and Sufi (2009) and Roberts (2015) shows 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 modify the baseline setting to provide a rationale for dynamically incomplete contracts with unspecified *interim* renegotiation.

Imagine players have some some idea about how much the asset will be worth two dates from now. Despite this, the players cannot elucidate the precise path the asset will take to evolve from what it is today to what the players expect it to be in two dates. Then, intuitively, the players will initially negotiate a dynamically incomplete contract that matures two dates from now, but contains missing details about what to do at date 1. Then, at date 1, once players observe the interim state of the world, they can use the continuation of the original contract as the point of departure for interim renegotiation. I now formalize this intuition.

Definition. *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 asset is left alone at date 1 - that is,* $k_1(s_1) = v_1$.

To capture both the extended horizon of the players and their uncertainty about the interim, I modify the baseline setting so that,

Assumption. $\mathcal{S}_1 = \{s_1\}$. s_0 *is a date 0 state of the world with the property that there is a set of probability distributions of* v_2, *call it* V_2, *such that* $\pi_1 \in \Pi_1$ *if and only if* $\Pi_2 \circ \pi_1 \subset V_2$.

Theorem 3. *Every Pareto-optimal payoff can alternatively be achieved by a dynamically incomplete contract with unspecified interim renegotiation.*

Proof. See Appendix. □

The meaning of this result mirrors that of Theorem 1. Fix a point on the Pareto-frontier. There are many renegotiation-proof contracts that achieve this payoff point. The entire collection is a union of sets, each of which can be characterized as follows: There exists a split $(\alpha^*_E(v_2), \alpha^*_F(v_2))$ of v_2 such that a contract belongs in the set if and only if it is renegotiation-proof at date 1 and every date 1 continuation contract Pareto-dominates leaving the asset alone at date 1 and then enacting the split - that is,

$$\begin{aligned}&\Big(u_E(c_{1,E}(v_1,x_2)) + \min_{\pi_2\in\Pi_2(k_1)} \mathbf{E}_{v_2\sim\pi_2} u_E(c_{2,E}(v_1,x_2,v_2)),\\ &\quad u_F(c_{1,F}(v_1,x_2)) + \min_{\pi_2\in\Pi_2(k_1)} \mathbf{E}_{v_2\sim\pi_2} u_F(c_{2,F}(v_1,x_2,v_2))\Big)\\ \geq \Big(&u_E(0) + \min_{\pi_2\in\Pi_2(v_1)} \mathbf{E}_{v_2\sim\pi_2} u_E(\alpha^*_E(v_2)), u_F(0) + \min_{\pi_2\in\Pi_2(v_1)} \mathbf{E}_{v_2\sim\pi_2} u_F(\alpha^*_F(v_2))\Big)\end{aligned}$$

for every (v_1, x_2). This characterization of the set admits the following interpretation:

Consider a tentative agreement to leave the asset alone at date 1 and split v_2 according to some $(\alpha_E(v_2), \alpha_F(v_2))$. For simplicity, I let $(\alpha_E(v_2), \alpha_F(v_2))$ refer to the entire agreement. This agreement $(\alpha_E(v_2), \alpha_F(v_2))$ represents the outside options for E and F at date 1. It does not necessarily equal actual date 1 and date 2 consumption. Instead, it serves as the individual rationality constraints for E and F at date 1 as they renegotiate a Pareto-optimal continuation contract. Let us refer to the agreement $(\alpha_E(v_2), \alpha_F(v_2))$ as a *dynamically incomplete contract* since there are no instructions for how it should be renegotiated at date 1.

Under expected utility, players cannot rank dynamically incomplete contracts. However, the characterization of the set of renegotiation-proof complete contracts that achieve a common payoff implies that under the current max-min setup, optimal dynamically incomplete contracting is well-defined. Given an arbitrary agreement $(\alpha_E(v_2), \alpha_F(v_2))$, so long as it is always Pareto-improved at date 1 to *some* Pareto-optimal continuation contract, the players can assign a value to it while still being agnostic about *which* Pareto-optimal date 1 continuation contract will be chosen.

Thus, to achieve a point on the complete contracts Pareto-frontier, it suffices for E and F to initially negotiate only an optimal dynamically incomplete contract $(\alpha^*_E(v_2), \alpha^*_F(v_2))$ that leaves details concerning any interim renegotiation completely unspecified. Then, once date 1 arrives, E and F can renegotiate freely to some Pareto-optimal date 1 continuation contract that Pareto-dominates the continuation of the original agreement.

Conclusion

In this paper, I developed a theory of optimal dynamically incomplete contracting. In contrast to previous theories, I did not place ad-hoc restrictions on the contracting space, and, instead, developed the theory from first-principles by extending a class of max-min preferences into the dynamic setting. I first established a "blank slate" optimality result which stated the optimality of dynamically incomplete contracts but made no predictions about the structure of the contract itself or how the missing details should be completed over time. Then I performed a series of modifications of the model - changes in the resolution of uncertainty and the observability of information - and investigated how they affected the structure of the initial dynamically incomplete contract, the scope of subsequent renegotiation, the necessity of control rights, and the allocation of control rights. From this analysis emerged a robust rationale for equity contracting and a micro-foundation for debt. Lastly, I provided a glimpse of the many possible future directions one can go from here.

5 Appendix

Proof of Lemma 1. Without loss of generality, I prove the result for E. Fix a stopping time $\tau \in \{1,2\}$ 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 $\tau(s_1) > 1$. By construction, the remaining capital in s_1 is the same under c' and c''. Call it $k_1(s_1)$. Then $U'_{1,E}(s_1) = u_E(c'_1(s_1)) + \min_{\pi_2\in\Pi_2(k_1(s_1))}\mathbf{E}_{s_2\sim\pi_2}U'_{2,E}(s_1,s_2) = u_E(c''_1(s_1)) + \min_{\pi_2\in\Pi_2(k_1(s_1))}\mathbf{E}_{s_2\sim\pi_2}U'_{2,E}(s_1,s_2) \leq u_E(c''_1(s_1)) + \min_{\pi_2\in\Pi_2(k_1(s_1))}\mathbf{E}_{s_2\sim\pi_2}U''_{2,E}(s_1,s_2) = U''_{1,E}(s_1)$. 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, obviously, $U'_{0,E} \leq U''_{0,E}$. $\square$

Proof of Theorem 2. For each date 1 state of the world s_1, define $\overline{U}_{1,F}(s_1)$ to be F's continuation payoff under 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 $s''_1 = (v''_1, \theta''_{1,E}, \Pi''_2)$ satisfying $U_{1,F}(s''_1) < \overline{U}_{1,F}(s''_1)$. Then for every $v'_1 \geq v''_1$, there exists a state $s'_1 = (v'_1, \theta'_{1,E}, \Pi'_2)$ 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 $s'_1 = (v'_1; \theta'_{1,E}, \Pi'_2 : (\hat{k}_1, \delta) \to 2^{\Delta(\{s_2\})})$ with the following properties: $\Pi'_2(k_1) = \Pi''_2([k_1 - (v'_1 - v''_1)] \wedge 0)$ and $\Pi'_2(\hat{k}_1, \delta) = \Pi'_2(k_1)$ for all $\delta \leq v'_1 - v''_1$.

Moreover, assume that $\theta'_{1,E}$ is small enough such that the continuation contract in state s'_1 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 $c_{1,F} < v''_1$ in state s'_1. To see why, first compute the continuation contract that delivers $\overline{U}_{1,F}(s''_1)$ to F in state s''_1. 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 all the consumption quantities for F and give them to E until the continuation payoff to F decreases to $U_{1,F}(s''_1)$. By construction of s'_1, this same continuation contract will also deliver $U_{1,F}(s''_1)$ to F in s'_1. Second, given that it is possible to deliver continuation payoff $U_{1,F}(s''_1)$ to F in state s'_1 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 in state s'_1 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 $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(s_2), c_{2,F}^*(s_2)\}$ be this continuation contract.

Suppose it is not true that $U_{1,F}(s_1') \leq U_{1,F}(s_1'')$. Then consider the case when 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 $U_{1,F}(s_1'')$ which, by assumption, is strictly smaller than what is promised to F if E tells the truth. Now consider the continuation contract $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(s_2), c_{2,F}^*(s_2)\}$. If E had told the truth, then F would value this continuation contract at $U_{1,F}(s_1'')$. However, because I assume that $\Pi_2'(\hat{k}_1, \delta) = \Pi_2'(k_1)$ 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 $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(s_2), c_{2,F}^*(s_2)\}$, E is made strictly better off and F thinks he is equally well-off. Now just tweak $\{c_{1,E}^*, c_{1,F}^*, c_{2,E}^*(s_2), c_{2,F}^*(s_2)\}$ 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 debt with the above D_1 to the Pareto-optimal contract.

Fix a distribution $\pi_1 \in \Pi_1$. 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'. Given the structure of the setting $(s_0, \mathcal{S}_1)$, $\pi_1' \in \Pi_1$.

Fix any state s_1 with capital value $v_1 < D_1$. By assumption, if $U_{1,F}(s_1) < \overline{U}_{1,F}(s_1)$ then $U_{1,F} \geq D_1 > v_1$. On the other hand, $\overline{U}_{1,F}(s_1) \geq v_1$. Thus, $U_{1,F}(s_1) \geq v_1$. If, furthermore, F's belief function is trivial, then $U_{1,F}(s_1) = \overline{U}_{1,F}(s_1) = v_1$. 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 debt weakly decreases moving from π_1 to π_1'. Moreover, the value of debt 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 debt 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 $\pi_1 \in \Pi_1$. 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'. $\pi_1' \in \Pi_1$. The value of the Pareto-optimal contract weakly decreases moving from π_1 to π_1'. Define a similar modified distribution π_1'' for debt. The value of debt weakly decreases moving from π_1 to π_1''. Moreover, the value of debt under π_1'' is weakly larger than the value of the Pareto-optimal contract under π_1'. Thus, E weakly prefers debt.

Step 3. Debt 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 misreport to are those where F retains control and E gets nothing. Thus, there is no way he can misreport to increase his payoff.

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 misreports 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 asset's productivity by misreporting. Again, there is no way he can misreport to a state and renegotiate the contract to make himself strictly better off. □

Proof of Theorem 3. Fix a contract (c_E^*, c_F^*) that achieves some Pareto-optimal payoff $(U_{0,E}^*, U_{0,F}^*)$. Fix an $\varepsilon > 0$ and define the date 1 state of the world

$$\hat{s}_1^\varepsilon = \left(\varepsilon; \hat{\Pi}_2 : k_1 \to 2^{\Delta(\{s_2\})}\right)$$

with the following properties: $\hat{\Pi}_2(\varepsilon) = V_2$, and $\hat{\Pi}_2(\delta) = 0$ for all $\delta < \varepsilon$. By construction, the belief $\hat{\pi}_1^\varepsilon \in \Delta(\mathcal{S}_1)$ that puts all weight on $\hat{s}_1^\varepsilon$ is in Π_1.

Consider the continuation contract following $\hat{s}_1^\varepsilon$. If there is any date 1 consumption by either player, then $(U_{1,E}(\hat{s}_1^\varepsilon), U_{1,F}(\hat{s}_1^\varepsilon)) \leq (u_E(\varepsilon), u_F(\varepsilon))$. This combined with the fact that $\hat{\pi}_1^\varepsilon \in \Pi_1$ means that if ε is sufficiently small, there will be no consumption at date 1 and $k_1(\hat{s}_1^\varepsilon) = \varepsilon$. In this case, define $\alpha_E^*(v_2) := c_{2,E}^*(\hat{s}_1^\varepsilon, v_2)$ and $\alpha_F^*(v_2) := c_{2,F}^*(\hat{s}_1^\varepsilon, v_2)$. Thus, the continuation payoffs satisfy

$$(U_{1,E}^*(\hat{s}_1^\varepsilon), U_{1,F}^*(\hat{s}_1^\varepsilon)) = \left(\min_{\pi_2 \in V_2} \mathbf{E}_{v_2 \sim \pi_2} u_E(\alpha_E^*(v_2)), \min_{\pi_2 \in V_2} \mathbf{E}_{v_2 \sim \pi_2} u_F(\alpha_F^*(v_2))\right).$$

Since $\hat{\pi}_1^\varepsilon$ is just one element of Π_1, it must be that

$$(U_{0,E}^*, U_{0,F}^*) \leq (U_1^*(\hat{s}_1^\varepsilon), U_1^*(\hat{s}_1^\varepsilon)).$$

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

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

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

Now the date 1 continuation payoff for each state s_1 is

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

Finally, to make the contract renegotiation-proof at date 1, weakly Pareto improve each continuation contract to a Pareto-optimal one. □

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] Antic, N. (2014) "Contracting with Unknown Technologies," working paper

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

[4] Che, Y.-K. and D. B. Hausch (1999) "Cooperative Investments and the Value of Contracting," *American Economic Review* Vol. 89, pp. 125-147

[5] Chung, T.-Y. (1991) "Incomplete Contracts, Specific Investments, and Risk Sharing," *Review of Economic Studies* Vol. 58, pp. 1031-1042

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

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

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

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

[10] Hebert, B. (2015) "Moral Hazard and the Optimality of Debt," working paper

[11] Innes, R. D. (1990) "Limited Liability and Incentive Compatibility with Ex-Post Information Asymmetry," *Journal of Economic Theory* Vol. 52, pp. 45 - 67

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

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

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

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

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

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

[18] Yang, M. (2015) "Optimality of Debt under Flexible Information Acquisition," working paper