Breakthroughs, Deadlines, and the Nature of Progress: Contracting for Multistage Projects

Brett Green
UC Berkeley, Haas

Curtis R. Taylor
Duke University

December 2015

Motivating example

Consider an entrepreneur (Anne) with an idea for a new product.

- Profiting from the idea requires completing two stages, e.g.,
  1. Develop a prototype
  2. Devise a cost-effective manufacturing process
- Each stage takes time, but the amount of time (and funding) needed is uncertain. Anne has no wealth.

A VC firm wants to fund the project, but faces several agency considerations:

- Anne can divert funds for private benefit, delaying progress
- Anne can (potentially) falsify or hide the prototype

Question: How should the VC structure the contract?

- How does it depend on Anne's ability to falsify/hide progress?

More generally

- Many projects require successful completion of multiple stages before their benefits can be realized
  - Venture capital
  - Product development
  - Basic research
- Incentives of entreprenuer/employee/scientist need not be aligned with funder/firm/institution
  - Moral hazard: Agent can shirk/divert resources
  - Private information: About progress made

This Paper: explore incentive provision for multistage projects.

- How does the nature of progress affect the optimal mechanism?

What is progress?

Limited value in isolation; but a necessary step or "piece of the puzzle"

- **Tangible**: easily observable, difficult to manipulate
  - License, permit or patent
  - FDA approval
  - Number of users

- **Intangible**: privately observed by the agent, easy to manipulate
  - An idea or approach
  - An algorithm or subroutine of a larger program
  - Independently conducted experimental results

Research questions

1. How to design incentives for multistage projects?

2. How does the "nature of the progress" affect the terms of financing?
   - If progress is tangible how is it used?
   - If intangible, is there a role for communication? Is it possible to elicit progress reports and use them in a meaningful way?

3. What are the implications for how to design projects?

Findings

1. How to design optimal incentives for multistage projects?
   - With multiple deadlines and a reward that decreases over time
2. How does the "nature of the progress" affect the optimal contract? If progress is tangible, how is it used?
   - When you make progress, it extends your clock, i.e., you get another round of funding.
   - If intangible, is it possible to elicit progress reports and use them in a meaningful way?
     - Yes, doing so optimally involves using a "soft deadline"
3. In designing projects and reporting requirements
   - There are benefits to
     - Imposing a "small" reporting cost (e.g., time, effort, paperwork)
     - Making the first stage somewhat "harder"
     - Eliciting information rather than obscuring it

Related theoretical literature

Dynamic agency

- Green (1987); Spear and Srivastava (1987); Phelan and Townsend (1991); Quadrini (2004); Clementi and Hopenhayn (2006); DeMarzo and Sannikov (2006); DeMarzo and Fishman (2007); Sannikov (2008)...

With (observable) Poisson arrivals

- Hopenhayn and Nicolini (1997); Biais et al. (2010); Hoffmann and Pfeil (2010); Piskorski and Tchistyi (2011); Mason and Välimäki (2011); DeMarzo et al. (2014)

Persistent private information

- Fernandes and Phelan (2000); Battaglini (2005); Zhang (2009); Williams (2011); Edmans et al. (2012); Hu (2014); Guo and Hörner (2015)

(Simultaneous) Multi-tasking

- Holmstrom and Milgrom (1991); Laux (2001); Varas (2014)

Empirical "support"

Tangible progress $\Rightarrow$ VC Contracts

- Staged financing contingent on observable milestones
- Entrepreneur's stake grows conditional on reaching milestones
- Right for VC to stop funding, control future financing, and terminate entrepreneur

These features arise as part of the optimal contract for a multistage project with tangible progress.

Anectdotal evidence

Intangible progress $\Rightarrow$ Grant funding for scientific research

- Extensive reporting requirements
- Indefinite penalties for lack of reported progress
  - According to the NSF Grant Policy Manual:

    NSF reserves the right...to withhold future funding after a specified date if the recipient fails to comply with the conditions of an NSF grant, including the reporting requirements.
  - Similarly, the NIH website states that:

    If your [progress] report is extremely late, you risk losing funding...

Consistent with (costly) reporting playing a crucial role and vagaries have "similar flavor" to soft deadline.

What this paper is not about

We do **not** study learning or experimentation, e.g.,

- Levitt and Snyder (1997), Inderst and Mueller (2010), Manso (2011), Bonatti and Hörner (2011,2014), Hörner and Samuelson (2014), Moroni (2015), Halac et al. (2015)

In our model, information about progress made (or lack thereof) has **no value** to a social planner

- Focus on using progress/reports to control agency costs

Model

A principal (P) contracts with an agent (A) to complete a project.

- Project requires: funding from P, expertise from A
  - P incurs flow cost to fund project prior to termination
- Completion of two stages or "breakthroughs" required to realize project benefits
- Arrival rate of breakthroughs depends on the A's hidden action
  - A can "shirk" (or divert investment) for private benefit
- A contract is a termination rule T and reward schedule R
  - Simple: single deadline and reward as function of success date
  - Complex: deadline contingent on history of progress or reports

Model details

- Both players risk neutral, A has limited liability (trivial otherwise)
- Continuous time, no discounting (for simplicity)
- P has full commitment power

Notation

- Benefit to P if project succeeds in ultimate stage: Π
- Flow cost of operation: c
- Agent action: $a_t \in \{0, 1\}$
  - Arrival rate of breakthrough: λa_t
  - Private flow benefit from diverting resources: ϕ
- Arrival time of breakthroughs: τ_1, τ_2

First best

Assumption

The project has positive expected value

$$\Pi - \frac{2c}{\lambda} > 0,$$

and shirking (or diversion) is inefficient

$$\phi \leq c.$$

Therefore, the first-best policy involves no shirking and no termination

- There is **no learning** about project quality in the model

The contract

P offers A a contract $\mathcal{C} = \{W, T\}$

- $dW_t \geq 0$: payment to agent at time t
  - Optimal backload all payments to agent (Ray, 2002)
- T: termination policy (or stopping rule)
  - Upon termination, players get their outside options (set to zero)

$\mathcal{C}$ must be measurable w.r.t. P's information set, which includes a randomization device.

Example

A **simple contract** consists of a single deadline, $T \in \mathbb{R}_+$, and a reward that depends only on the time of success.

Payoffs

- To the principal:

$$\mathcal{P}_0(\mathcal{C}) = \mathbb{E}^a \left[\Pi \cdot \mathbb{1}_{\{\tau_2 \leq T\}} - \int_0^{T \wedge \tau_2} (cdt + dW_t) \right]$$

- To the agent:

$$U_0(\mathcal{C}) = \mathbb{E}^a \left[\int_0^{T \wedge \tau_2} \left(\phi \mathbb{1}_{\{a_t = 0\}} dt + dW_t \right) \right]$$

(P-)Optimal contract: maximizes $\mathcal{P}_0(\mathcal{C})$ subject to relevant constraints

The nature of progress

We focus on the following two extreme cases:

Definition

- *Progress is* **tangible** *if* τ_1 *is publicly observable and verifiable*
- *Progress is* **intangible** *if* τ_1 *is privately observed by* A *and is not verifiable*

Throughout, the ultimate success of the project (i.e., τ_2) is tangible.

- Observability not crucial, verifiability is.

One-stage benchmark

Suppose that only one breakthrough is needed to realize Π

- Denote the reward for success at t by R_t
- U_t is the agent's continuation value at t

Preliminaries

1. It is optimal for the agent not to shirk/divert iff

$$\lambda(R_t - U_t) \geq \phi \tag{IC}$$

2. It is WLOG to focus on $\mathcal{C}$ with no shirking

One-stage problem

Via standard recursive methods (Spear and Srivastava, 1987)

- State variable is A's continuation value: u

The principal's value function solves the HJB equation

$$\lambda V(u) = \max_R \lambda(\Pi - R) - c + V'(u)\frac{du}{dt}$$

subject to $V(0) = 0$ and

$$\lambda(R - u) \geq \phi \tag{IC}$$

$$\frac{du}{dt} = -\lambda(R - u) \tag{PK}$$

Solution to the one-stage problem

Result

In a one-stage project, (IC) always binds. Therefore,

$$\frac{du}{dt} = -\phi \quad \text{and} \quad R(u) = u + \frac{\phi}{\lambda}$$

and the principal's value function is given by

$$V_{\text{one stage}}(u) = \left(\Pi - \frac{c}{\lambda}\right)\left(1 - e^{-\lambda u/\phi}\right) - u$$

One-stage implementation

Result

The P-optimal contract for a one-stage project can be implemented with a simple contract, where the deadline is

$$T^* = \frac{1}{\lambda} \ln \left(\frac{\lambda \Pi - c}{\phi} \right)$$

and the reward schedule decreases linearly over time

$$R_t = \frac{\phi}{\lambda} + \phi(T^* - t).$$

The optimal deadline is decreasing in the severity of agency conflict

$$\lim_{\phi \to 0} T^* = \infty$$

.

One-stage payoff frontier

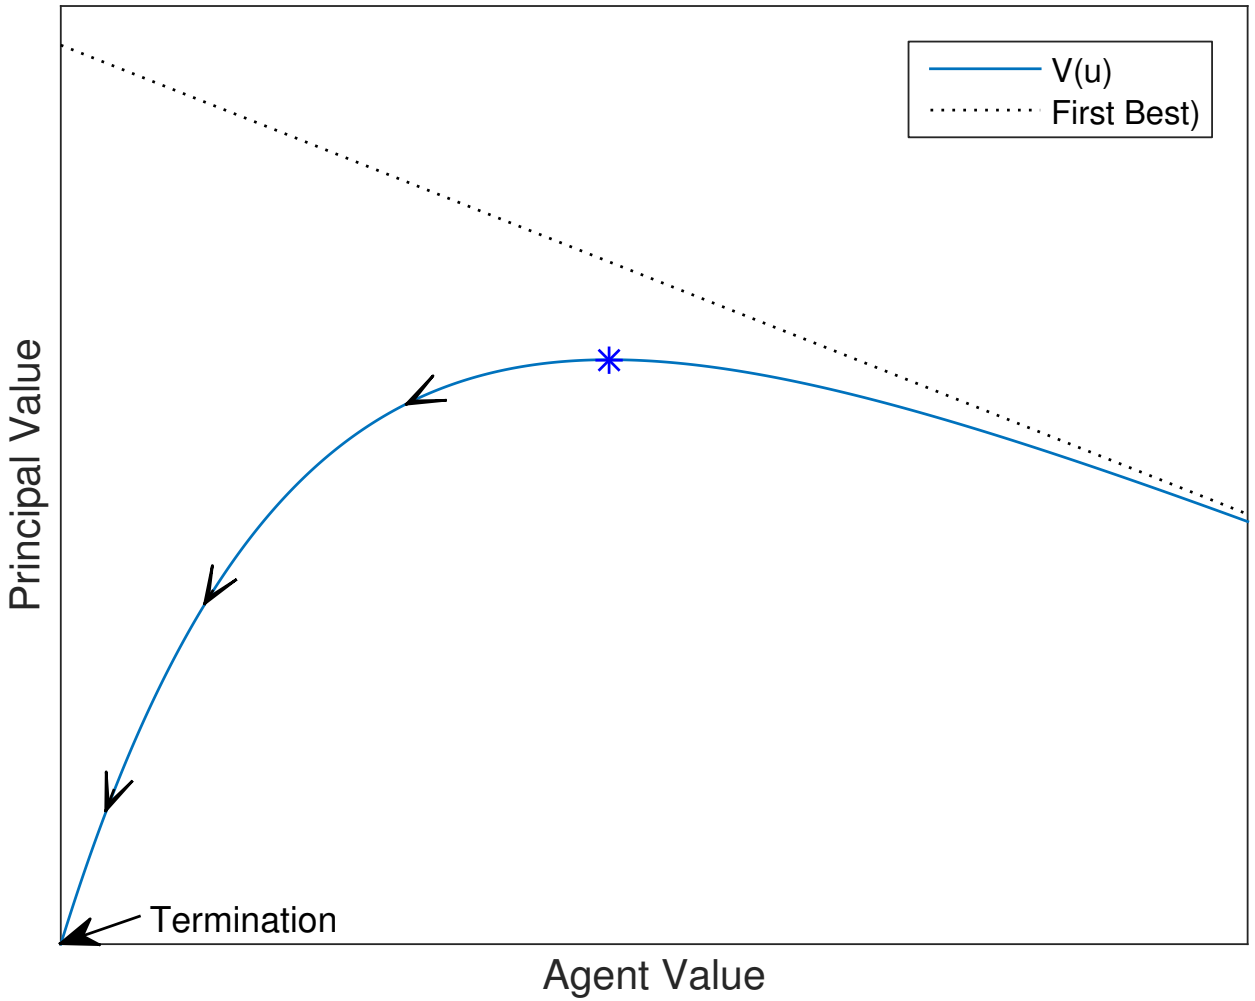

Agent continuation value decreases over time as deadline approaches.

Tangible progress

Solve by backward induction on the project stage

- State variable is a pair $(s, u) \in \{1, 2\} \times \mathbb{R}_+$
- $V_s(u)$ is P's value function in stage s
- Characterize optimum through dynamics of A's continuation value

Result

The optimal contract involves no shirking. The incentive compatibility constraint in stage s requires that

$$\lambda(R_s(u) - u) \geq \phi, \tag{IC}$$

where $R_s(u)$ is the promised utility to the agent for a success in stage s given u. Moreover, (IC) is binding for all u in the second stage.

Second-stage payoffs = One-stage benchmark

Given u, we already solved the second-stage problem!

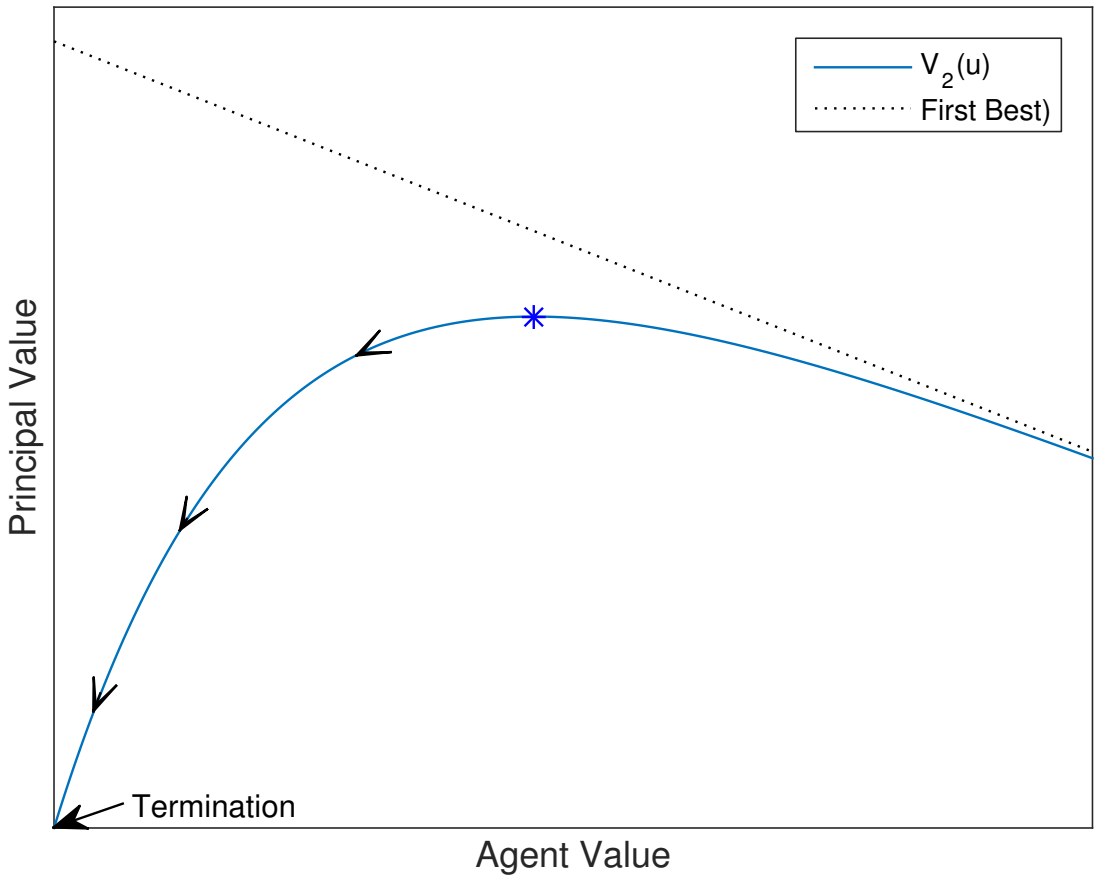

Agent continuation value decreases over time as deadline approaches.

Solution to the first stage problem

Two cases:

1. IC always binds in the first stage
   - Occurs when ultimate success is sufficiently valuable

$$\Pi \geq \frac{2c}{\lambda}\left(\frac{e-1}{e-2}\right) \tag{C1}$$

2. IC is slack for some u in the first stage
   - Occurs when (C1) does not hold
   - Optimal to set $R_1(u) = \underline{w}$ for all $u \in [0, \underline{u}]$
   - P optimally shifts resources from first to second stage

When (C1) holds

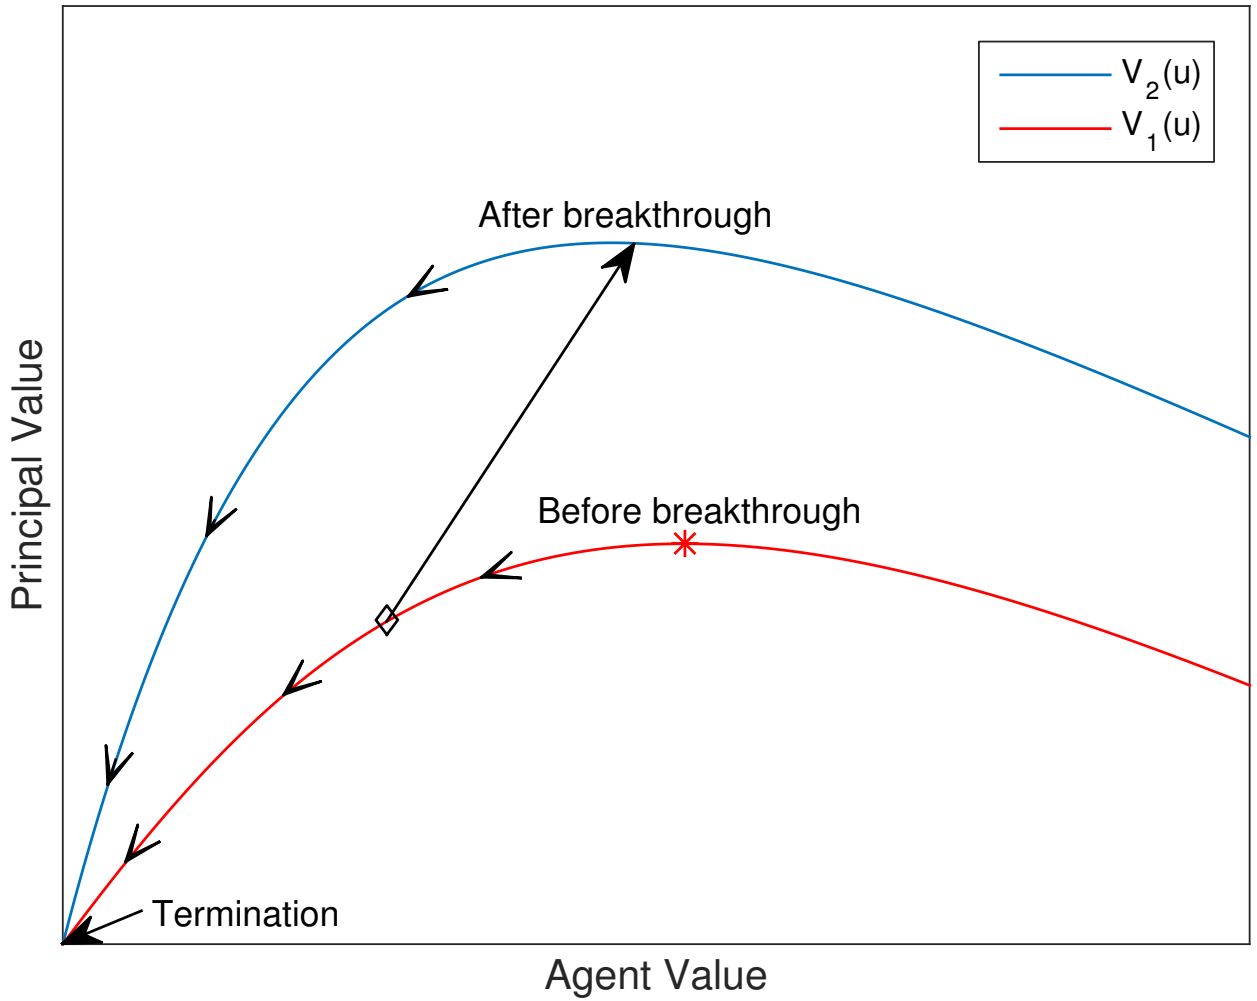

Continuation value jumps up by ϕ/λ following a breakthrough regardless of when it occurs

When (C1) does not hold

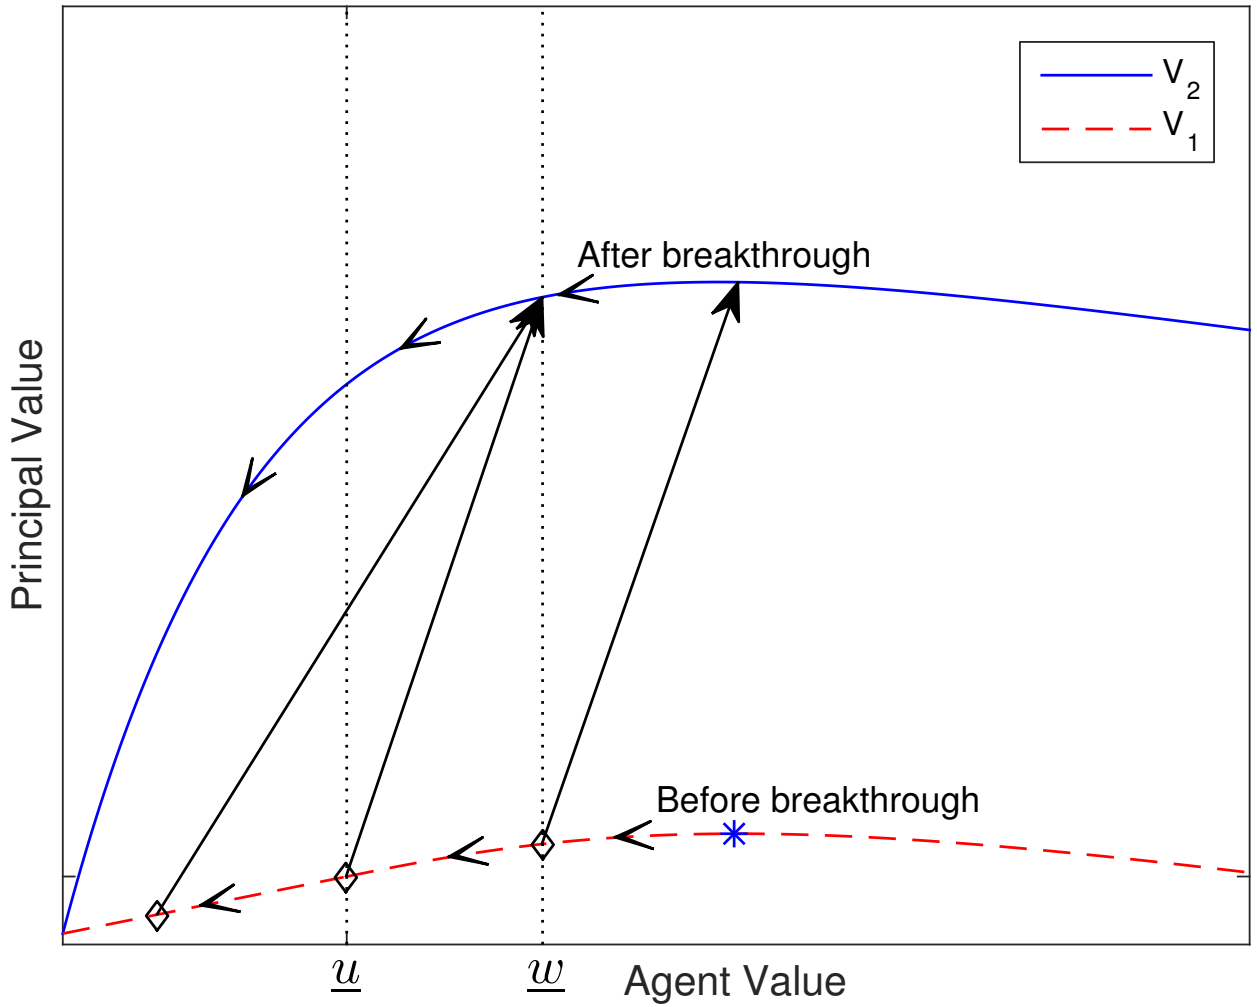

Continuation value increases by more than ϕ/λ if $u < \bar{u}$

Optimal termination policy

When (C1) holds, principal uses two "hard" deadlines, $T_1 < T_2$.

- If no breakthrough before T_1
  - Project terminated
- If first breakthrough before T_1, A gets until T_2 to make second breakthrough
  - Agent rewarded only if $\tau_2 \leq T_2$

Equivalent to one deadline with fixed "extension" for success

When (C1) does not hold, the second deadline depends on τ_1

- If no breakthrough by t_s, the second clock is "paused" and restarted only after a breakthrough

Equivalent to one deadline with "extension" that depends on τ_1

Intangible progress

- The agent privately observes the first breakthrough
- The second breakthrough is still observable/contractible

A few observations:

1. Agent acquires (persistent) private information over time
   - Call the agent types A_1 (low) and A_2 (high)
2. The optimal contract with tangible progress would not induce truth telling
   - A_1 would falsely report breakthrough just before T_1
   - A_2 would hide a breakthrough and shirk for $t \in (t_s, T_1]$
3. A simple contract (independent of reports) would trivially induce truth telling..
   - But can the principal do better?

Without communication

Lemma

For any contract with fixed deadline T and bounded reward scheme, there exists a $\bar{t}$ such that A_1 will shirk for all $t \in (\bar{t}, T]$.

Why?

- Probability of two breakthroughs proportional to $(\lambda dt)^2$
- Shirking yields payoff proportional to dt

But shirking is an inefficient form of compensation $(\phi \leq c)$

- Might as well just pay A_1 and terminate the project.

However, a severance payment is also not optimal...

With communication

By the revelation principle, can restrict attention to direct mechanisms in which the agent to reports progress truthfully and immediately.

- The truthtelling constraints take two forms

1. No False Progress (NFP): Agent does not want to falsely report a breakthrough.
2. No Hidden Progress (NHP): Agent does not want to "hide" a breakthrough.
    - Why hide progress? Think grant proposal...

Intangible progress $\Longrightarrow$ persistent private information

To formulate principal's problem recursively, need three state variables

- s: Project stage (as reported by the agent)
- u_1: Promised utility to A_1
- u_2: Promised utility to A_2

With the extra state variable, we can take a similar approach to the case of tangible progress

- Solve the principal's problem by backward induction on s for any (u_1, u_2).

Implementable utility pairs

A_2 can always do at least as well as A_1

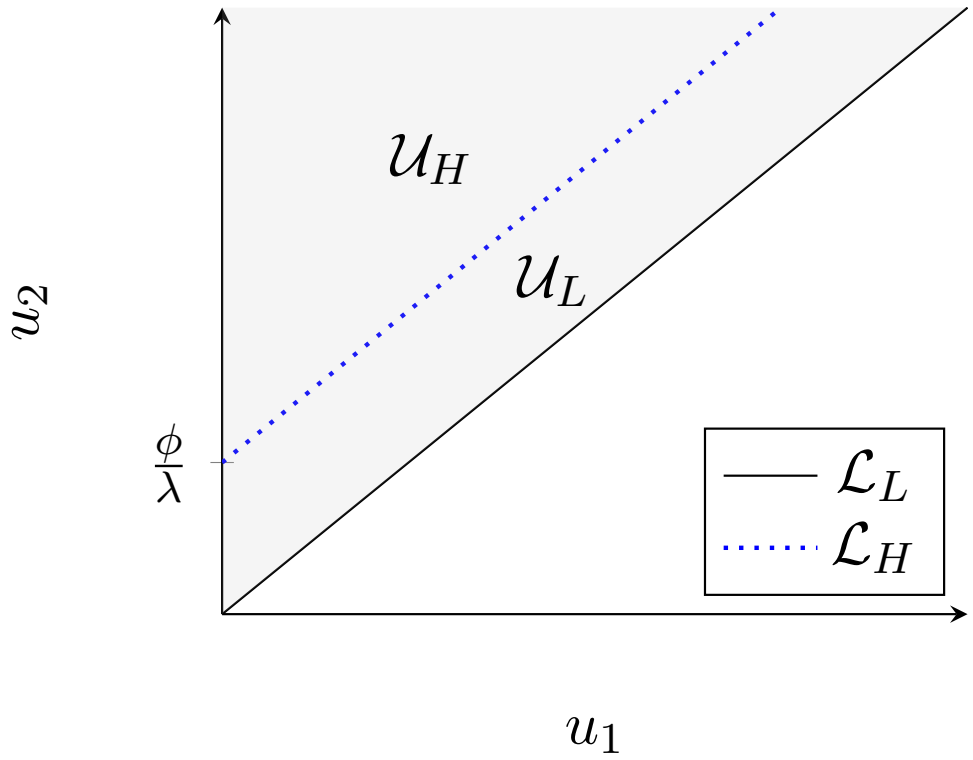

- Note that u_2 is the "reward" to A_1 for a breakthrough
- Hence, A_1 will shirk if state evolves in $\mathcal{U}_L$

After the first (reported) breakthrough

- Need to deliver promised utility u_2 to A_2 (PK)
- A_1 can get **no more than** u_1 (NFP)

Lemma

For any $u \in \mathcal{U}$, the principal's value function in the second stage is

$$F_2(u_1, u_2) = \underbrace{\left(1 - e^{-\lambda \frac{u_1}{\phi}}\right)}_{\mathbb{P}\left(\tau_2 \leq \frac{u_1}{\phi}\right)} \left(b - \frac{c}{\lambda}\right) - u_2$$

- Principal sets deadline of length $\frac{u_1}{\phi}$
  - Unlike with tangible progress, deadline is independent of u_2!
- R satisfies (IC) and (PK) for A_2

Optimal dynamics prior to breakthrough

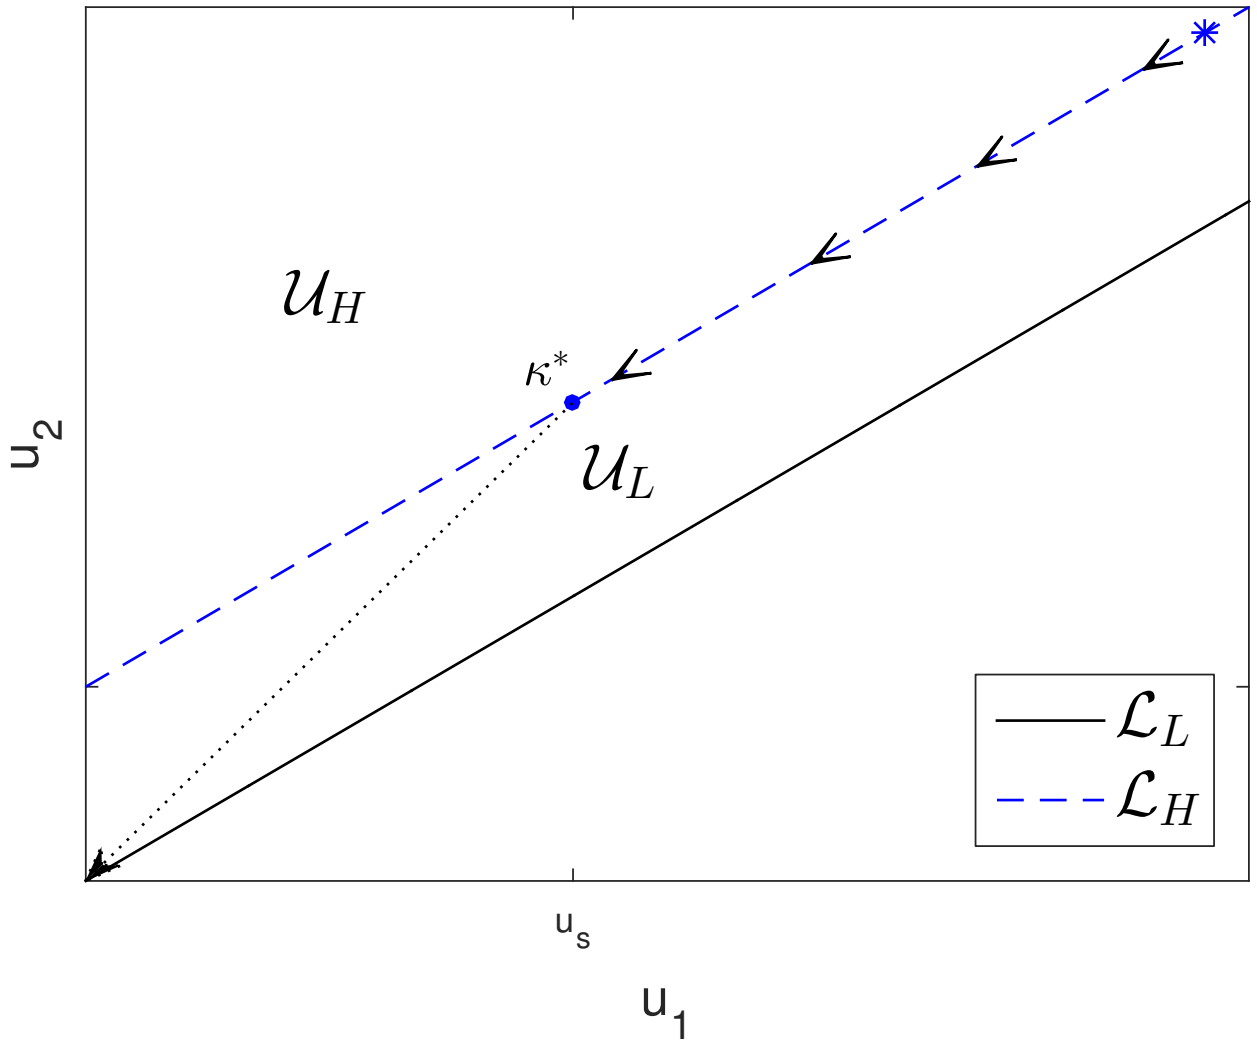

Continuation utilities drift down toward κ^*, communication prior to κ^* is not necessary.

What happens at κ^*?

Upon reaching κ^*, P asks: *"Have you made a breakthrough yet?"*

- If agent answers "yes"
  - Principal gives agent a relatively short deadline to make the final breakthrough.
- If agent answer "no"
  - Principal begins randomly terminating the project
  - Conditional on not terminating, state remains at κ^*

Interpretation: A "soft" deadline

- P guarantees funding until κ^*, but reserves the right to terminate thereafter if progress has not been reported
- Immediate communication is **critical** after the soft deadline

Intuition

Why is a soft deadline part of the optimal contract?

1. Why is it incentive compatible at κ^*?
   - A_2 likely to have ultimate success soon, does not want to risk termination
     $\implies$ strictly prefers truthtelling over the lottery (NHP slack).
   - A_1 is unlikely to succeed in near future, could "lie and steal" but weakly prefers the lottery by construction (NFP binds).

2. Why does the principal want to use it?
   - Needs to give A_1 utility of u_s (PK) **without increasing** CV after a breakthrough (NFP).
   - But P would like to give **more time** conditional on a breakthrough
   - Optimal to kill project today (with some prob) to preserve time following a breakthrough

Tangible vs Intangible

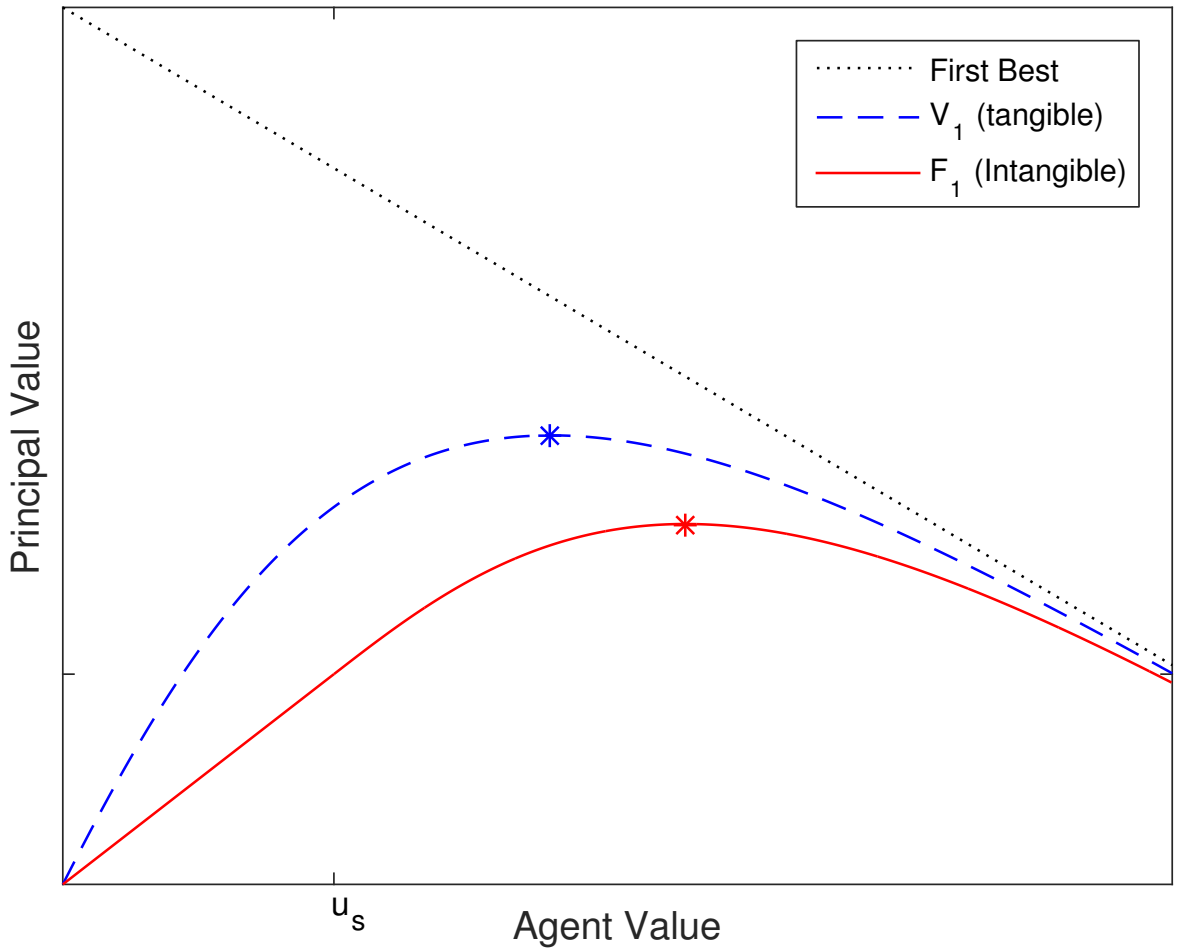

- Tangible: Principal does better, total surplus is higher
- Intangible: Agent earns more rents under P-optimal contract

Asymmetric stages

In many relevant applications, one stage may be

- Expected to take more time (smaller λ)
- Require more working capital (higher c)
- More susceptible to diversion/shirking (higher ϕ)

To fix ideas, keep c and ϕ constant across stages and parameterize the asymmetry by $\alpha \in [-1, 1]$, where

$$\frac{1}{\lambda_1} = \frac{1+\alpha}{\lambda} \quad \text{and} \quad \frac{1}{\lambda_2} = \frac{1-\alpha}{\lambda}$$

- Maintains project value at $\Pi - 2c/\lambda$ as α varies

Asymmetric stages

Result

With either tangible or intangible progress, the principal's ex-ante payoff under the P-optimal contract is

- Increasing in α for α near 0
- Decreasing in α for α near 1
- Converges to one-stage project value as $\alpha \to 1$

Intuition

- On the margin, it is cheaper to compensate agent with promised utility in stage 2 rather than cash for ultimate breakthrough
- But as $\alpha \to 1$, monitoring technology becomes inferior

Implication: Optimal to make the first stage more difficult than the second stage (but not too much more difficult).

Costly reporting

Formal channels of communication often require the agent incur costs to documenting progress.

Question: Can the principal benefit by imposing reporting costs?

- Benefit: relaxes (NFP), can give more time after a positive report
- Cost: requires extra compensation to A_2

Result

Let ρ denote the cost incurred by the agent to report progress.

- A direct mechanism with $\rho > 0$ is inferior to one with $\rho = 0$.
- For ρ small, there exists an (indirect) mechanism in which P's payoff is higher than the optimal mechanism with $\rho = 0$.

Intuition: Only require reports after κ^*

Information suppression?

With intangible progress, P must give A rents to induce truth telling

- But P also uses this information to optimize termination policy

Question: Do these benefits outweigh the costs?

- If P cannot observe progress, restrict A's ability to do so?

Result

When τ_1 is unobservable to both P and A:

- The P-optimal contract can be implemented with a simple contract
- The principal is worse off than with intangible progress

Implication: Better to elicit information than try to suppress it

Conclusion

Studied provision of incentives in multistage projects

- Nature of progress is important
  - Tangible: Sequence of hard deadlines resembling staged financing
  - Intangible: Communication is valuable, optimal to use "soft" deadlines
- Implications for project design
  - Better to make first stage slightly harder
  - Costly reporting can be valuable
  - Better to elicit rather than suppress

More progress yet to be made...

- Commitment, replacement, competition, continuous progress...
- Assuming NHP not violated :)

Battaglini, M. (2005). Long-term contracting with Markovian consumers. *American Economic Review 95*, 637–658.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risk, limited liability and dynamic moral hazard. *Econometrica 78*, 73–118.

Clementi, G. L. and H. A. Hopenhayn (2006). A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics 121*, 229–265.

DeMarzo, P. M. and M. J. Fishman (2007). Optimal long-term financial contracting. *Review of Financial Studies 20*, 2079–2128.

DeMarzo, P. M., D. Livdan, and A. Tchistyi (2014). Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives. *Working Paper*.

DeMarzo, P. M. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance 61*, 2681–2724.

Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov (2012). Dynamic CEO Compensation. *Journal of Finance 67*(5), 1603–1647.

Fernandes, A. and C. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory 91*, 223–247.

Green, E. J. (1987). Lending and the Smoothing of Uninsurable Income. In E. C. P. Wallace and Neil (Eds.), *Minnesota Studies in Macroeconomics, Volume I*. University of Minnesota Press.

Guo, Y. and J. Hörner (2015). Dynamic Mechanisms without Money. *Working Paper*.

Halac, M., N. Kartik, and Q. Liu (2015). Optimal Contracts for Experimentation. *Working Paper*.

Hoffmann, F. and S. Pfeil (2010). Reward for luck in a dynamic agency model. *Review of Financial Studies 23*, 3329–3345.

Holmstrom, B. and P. Milgrom (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics and Organization 7*, 24–52.

Hopenhayn, H. A. and J. P. Nicolini (1997). Optimal Unemployment Insurance. *Journal of Political Economy 105*, 412.

Hörner, J. and L. Samuelson (2014). Incentives for Experimenting Agents. *RAND Journal of Economics*.

Hu, Z. (2014). Financing Innovation with Unobserved Progress. *Working Paper*.

Inderst, R. and H. M. Mueller (2010). CEO replacement under private information. *Review of Financial Studies 23*, 2935–2969.

Laux, C. (2001). Limited-liability and incentive contracting with multiple projects. *RAND Journal of Economics 32*, 514–526.

Levitt, S. D. and C. M. Snyder (1997). Is no news bad news? Information transmission and the role of 'early warning' in the principal-agent model. *RAND Journal of Economics 28*, 641–661.

Manso, G. (2011). Motivating Innovation. *Journal of Finance 66*, 1823–1860.

Mason, R. and J. Välimäki (2011). Dynamic moral hazard and stopping. *. . . of Exeter and . . . .*

Moroni, S. (2015). Experimentation in Organizations. *Working Paper*.

Phelan, C. and R. M. Townsend (1991). Computing Multi-Period, Information-Constrained Optima. *Review of Economic Studies 58*, 853.

Piskorski, T. and A. Tchistyi (2011). Stochastic house appreciation and optimal mortgage lending. *Review of Financial Studies 24*, 1407–1446.

Quadrini, V. (2004). Investment and liquidation in renegotiation-proof contracts with moral hazard. *Journal of Monetary Economics 51*, 713–751.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies 75*, 957–984.

Spear, S. E. and S. Srivastava (1987). On Repeated Moral Hazard with Discounting. *The Review of Economic Studies 54*, 599.

Varas, F. (2014). Contracting Timely Delivery with Hard to Verify Quality. *Working Paper*.

Williams, N. (2011). Persistent Private Information. *Econometrica 79*(4), 1233–1275.

Zhang, Y. (2009). Dynamic contracting with persistent shocks. *Journal of Economic Theory 144*, 635–675.