



Cowles Foundation

for Research in Economics
at Yale University

Cowles Foundation Discussion Paper No. 1726RR

INCENTIVES FOR EXPERIMENTING AGENTS

Johannes Hörner and Larry Samuelson

September 2009
Revised March 2013

An author index to the working papers in the
Cowles Foundation Discussion Paper Series is located at:
<http://cowles.econ.yale.edu/P/au/index.htm>

This paper can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=2235747>

Incentives for Experimenting Agents*

Johannes Hörner
Department of Economics
Yale University
New Haven, CT 06520
Johannes.Horner@yale.edu

Larry Samuelson
Department of Economics
Yale University
New Haven, CT 06520
Larry.Samuelson@yale.edu

March 12, 2013

Abstract. We examine a repeated interaction between an agent, who undertakes experiments, and a principal who provides the requisite funding for these experiments. The repeated interaction gives rise to a dynamic agency cost—the more lucrative is the agent’s stream of future rents following a failure, the more costly are current incentives for the agent, giving the principal an incentive to reduce the continuation value of the project. We characterize the set of recursive Markov equilibria. We also find that there are non-Markov equilibria that make the principal better off than the recursive Markov equilibrium, and that may make both agents better off. Efficient equilibria front-load the agent’s effort, inducing as much experimentation as possible over an initial period, until making a switch to the worst possible continuation equilibrium. The initial phase concentrates the agent’s effort near the beginning of the project, where it is most valuable, while the eventual switch to the worst continuation equilibrium attenuates the dynamic agency cost.

*We thank Dirk Bergemann for helpful discussions and the editor and three referees for helpful comments. We thank the National Science Foundation (SES-0549946, SES-0850263 and SES-1153893) for financial support.

Contents

1	Introduction	1
1.1	Experimentation and Agency	1
1.2	Optimal Incentives: A Preview of Our Results	2
1.3	Applications	5
1.4	Related Literature	7
2	The Model	7
2.1	The Agency Relationship	7
2.1.1	Actions	7
2.1.2	Strategies and Equilibrium	9
2.2	What is Wrong with Markov Equilibrium?	10
2.3	The First-Best Policy	12
3	Characterization of Equilibria	13
3.1	Recursive Markov Equilibria	13
3.1.1	No Delay	13
3.1.2	Delay	19
3.1.3	Recursive Markov Equilibrium Outcomes: Summary	21
3.1.4	Recursive Markov Equilibrium Outcomes: Lessons	23
3.2	Non-Markov Equilibria	24
3.2.1	Lowest Equilibrium Payoffs	25
3.2.2	The Entire Set of (Limit) Equilibrium Payoffs	30
3.2.3	Non-Markov Equilibria: Lessons	32
3.3	Summary	33
4	Comparisons	33
4.1	Observable Effort	34
4.1.1	Markov Equilibria	34
4.1.2	Non-Markov Equilibria	36
4.1.3	Comparison	37
4.2	Good Projects: $\bar{q} = 1$	37
4.3	Commitment	39
4.4	Powerless Principals	39
5	Summary	41
A	Online Appendix: Foundations	A-1
A.1	Outline	A-1
A.2	Strategies	A-2
A.3	Expectations and the Agent's Incentives	A-2
A.4	Preliminaries	A-4
A.4.1	The Horizon is Effectively Finite	A-4
A.4.2	Two Beliefs are Enough	A-5
A.4.3	The Value of an Optimistic Agent	A-5
A.4.4	The Value of a Pessimistic Agent	A-6
A.5	The Equilibrium Concept: Recursive Markov Equilibria	A-6
A.6	A Candidate Markov Equilibrium: No Delay Principal Optimum	A-8

A.6.1	The Strategies	A-9
A.6.2	The Costs of Agency	A-9
A.6.3	Positive Principal Payoffs?	A-10
A.6.4	The Agent’s Incentives	A-12
A.6.5	Completing the Strategies	A-12
A.6.6	Summary: No-Delay Principal-Optimum Markov Equilibrium	A-15
A.7	Other No-Delay Markov Equilibria	A-15
A.7.1	The Final Period	A-15
A.7.2	Constructing the Set of No-Delay Equilibria	A-16
A.8	The Set of Markov Equilibria	A-16
A.8.1	A Canonical Equilibrium	A-16
A.8.2	Characterizing the Canonical Equilibrium	A-17
A.8.3	Summary: Canonical Markov Equilibrium	A-20
A.8.4	Limit Uniqueness	A-21
A.9	s^W vs. s^S , Agent Offers	A-21
A.10	Proof of Lemma 4	A-23
A.11	Proof of Lemma 5	A-24
A.11.1	The Final Period	A-24
A.11.2	The Penultimate Period	A-25
A.11.3	The Induction Step	A-27
A.12	Proof of Lemma 6	A-31
A.13	Proof of Lemma 7	A-32
A.14	Proof of Lemma 9	A-35
A.15	Proof of Lemma 10	A-38
A.16	Proof of Lemma 11: Agent	A-42
A.17	Proof of Lemma 11: Principal	A-42
A.17.1	Outline	A-42
A.17.2	The Value of Merging	A-43
A.17.3	The Relative Value of Merging: Intuition	A-45
A.17.4	The Relative Value of Merging: Formal Analysis	A-49
A.18	Proof of Lemma 12	A-53
A.19	Proof of Lemma 13	A-55
A.20	Proof of Lemma 14	A-57
A.21	Proof of Lemma 15	A-59
A.21.1	The No-Delay Principal-Optimum Markov Equilibrium	A-60
A.21.2	The Bound	A-61
A.22	Proof of Lemma 16	A-67
A.23	Proof of Lemma 18	A-69

B Online Appendix: Comparative Statics, The First-Best Policy **B-71**

C Online Appendix: Proofs, Unobservable Effort **C-72**

C.1	Proof of Lemma 3	C-72
C.1.1	The Agent’s Highest No-Delay Payoff	C-72
C.1.2	The Principal’s Lowest Payoff	C-73
C.1.3	The Agent’s Lowest Payoff	C-73
C.2	Proof of Proposition 2	C-75
C.2.1	A Preliminary Inequality	C-76

C.2.2	Front-Loading Effort	C-78
D	Online Appendix: Proofs, Observable Effort	D-79
D.1	Proof of Proposition 4	D-79
D.2	Proof of Proposition 5	D-82
E	Online Appendix: Derivations and Proofs, Good Projects	E-83
E.1	The First-Best Policy	E-83
E.2	Stationary No-Delay Equilibrium: Impatient Projects	E-84
E.3	Markov Equilibria for Other Parameters	E-85
E.3.1	Patient Projects: Delay	E-85
E.4	Non-Markov Equilibria	E-86
E.4.1	Impatient Projects	E-86
E.4.2	Patient Projects	E-91
E.4.3	Characterization of Equilibrium Payoffs	E-92
E.4.4	Summary	E-94
E.5	Proofs	E-95
E.5.1	Proof of Lemma 19	E-95
E.5.2	Proof of Lemma 20	E-96
E.5.3	Proof of Lemma 21	E-98
E.5.4	Proof of Lemma 22	E-99
E.5.5	Proof of Proposition 9	E-101
F	Online Appendix: Proofs, Commitment	F-102
F.1	Proof of Proposition 6	F-102
F.1.1	A Preliminary Inequality	F-102
F.1.2	Front-Loading Effort	F-103

Incentives for Experimenting Agents

Johannes Hörner and Larry Samuelson

1 Introduction

1.1 Experimentation and Agency

Suppose an agent has a project whose profitability can be investigated and potentially realized only through a series of costly experiments. For an agent with sufficient financial resources, the result is a conceptually straightforward programming problem. He funds a succession of experiments until either realizing a successful outcome or becoming sufficiently pessimistic as to make further experimentation unprofitable. But what if he lacks the resources to support such a research program, and must instead seek funding from a principal? What constraints does the need for outside funding place on the experimentation process? What is the nature of the contract between the principal and agent?

This paper addresses these questions. In the absence of any contractual difficulties, the problem is still relatively straightforward. Suppose, however, that the experimentation requires costly effort on the part of the agent that the principal cannot monitor (and cannot undertake herself). It may require hard work to develop either a new battery or a new pop act, and the principal may be able to verify whether the agent has been successful (presumably because people are scrambling to buy the resulting batteries or music), but unable to discern whether a string of failures represents the unlucky outcomes of earnest experimentation or the product of too much time spent playing computer games. We now have an incentive problem that significantly complicates the relationship. In particular, the agent continually faces the temptation to eschew the costly effort and pocket the funding provided for experimentation, explaining the resulting failure as an unlucky draw from a good-faith effort, and hence must receive sufficient rent to forestall this possibility.

The problem of providing incentives for the agent to exert effort is complicated by the assumption that the principal cannot commit to future contract terms. Perhaps paradoxically, one of the advantages to the agent of a failure is that the agent may then be able to extract further rent from future experiments, while a success obviates the need for the agent and terminates the rent stream. A principal with commitment power could reduce the cost of current incentives by committing to a string of less lucrative future contracts (perhaps terminating experimentation altogether) in the event of failure. Our principal, in contrast, can alter future contract terms or terminate the relationship only if doing so is sequentially rational.

1.2 Optimal Incentives: A Preview of Our Results

The contribution of this paper is threefold.

First, we make a methodological contribution. Because the action of the agent is hidden, his private belief may differ from the public belief held by the principal. We develop techniques to solve for the equilibria of this hidden-action hidden-information problem. We work with a continuous-time model which, in order to be well defined, incorporates some inertia in actions, in the form of a minimum length of time dt between offers on the part of the principal.¹ The bulk of the analysis, beginning with Appendix A (which is posted online, as are all other appendices), provides a complete characterization of the set of equilibria for this game, and it is here that we make our methodological contribution. However, because this material is detailed and technical, the paper (in Sections 3–4) examines equilibrium outcomes in the frictionless limit obtained by letting dt go to zero. These arguments are intuitive and may be the only portion of the analysis of interest to many readers. One must bear in mind, however, that this is not an analysis of a game without inertia, but a description of the limits (as $dt \rightarrow 0$) of equilibrium outcomes in games with inertia. Toward this end, the intuitive arguments made in Sections 3–4 are founded on precise arguments and limiting results presented in Appendix A. Solving the model also requires formulating an appropriate solution concept in the spirit of Markov equilibrium. As explained in Section 3.1, to ensure existence of equilibrium, we must consider a weakening of Markov equilibrium, which we refer to as recursive Markov equilibrium.

Second, we study the optimal provision of incentives in a dynamic agency problem. As in the static case, the provision of incentives requires that the agent be given rents, and this agency cost forces the early termination of the project. More interestingly, because a success also terminates the project, the agent will exert effort only if compensated for the potential loss of his continuation payoff that higher effort makes more likely. This gives rise to a dynamic agency cost that shapes the structure of equilibrium. Compensating the agent for this opportunity cost can be so expensive that the project is no longer profitable to the principal. The project must be then downsized (formally, the rate of experimentation must be slowed down), to bring this opportunity cost down to a level that is consistent with the principal breaking even. Keeping the relationship alive can thus call for a reduction in its value.

When does downsizing occur? It depends on conflicting forces. The faster the rate at which learning occurs, or the lower the common prior belief, the shorter the maximum possible duration of experimentation before the project is abandoned for good. This brings down the continuation payoff of the agent, and so helps prevent the occurrence of

¹There are well-known difficulties in defining games in continuous time, especially when attention is not restricted to Markov strategies. See, in particular, Bergin and MacLeod [4] and Simon and Stinchcombe [28]. Our reliance on an inertial interval between offers is similar to the approach of Bergin and MacLeod.

downsizing. On the other hand, a lower common prior belief also means lower expected profitability, making it harder for the principal to break even without downsizing. As a result, whether and when downsizing occurs depends on the players' beliefs, the players' patience, and the rate of learning. While the (limiting) recursive Markov equilibrium outcome is always unique, downsizing in this equilibrium might occur either for beliefs below, or for beliefs above a given threshold. If a failure conveys a substantial amount of information, for example, then it must be that a success is likely *if* the project is good. Hence, exerting effort is risky for the agent, especially if he thinks it quite likely that the project is good. As a result, a high rate of learning leads to downsizing for high prior beliefs of the principal, who would be better off if she were more pessimistic! We might then expect optimistic agents who anticipate learning about their project quickly to rely more heavily on self-financing or debt contracts, not just because their optimism portends a higher expected payoff, but because they face particularly severe agency problems.

Recursive Markov equilibria highlight the role of the dynamic agency cost. Because their outcome is unique, and the related literature has focused on similar Markovian solution concepts, they also allow for clear-cut comparisons with other benchmarks discussed below. However, we believe that non-Markov equilibria better reflect (for example) actual venture capital contracts (see Section 1.3 below). The analysis of non-Markov equilibria is carried out in Section 3.2, and builds heavily on that of recursive Markov equilibria. Unlike recursive Markov equilibria, constrained efficient (non-Markov) equilibria always have a very simple structure: they feature an initial period where the project is operated at maximum scale, before the project is either terminated or downsized as much as possible (given that, in the absence of commitment, abandoning the project altogether need not be credible). So, unlike in recursive Markov equilibria, the agent's effort is always front-loaded. The principal's preferences are clear when dealing with non-Markov equilibria—she always prefers a higher likelihood that the project is good, eliminating the non-monotonicity of the Markov case. The principal always reaps a higher payoff from the best non-Markov equilibrium than from the Markov equilibrium, and the non-Markov equilibria may make both players better off than the Markov equilibrium.

It is not too surprising that the principal can gain from a non-Markov equilibrium. Front-loading effort on the strength of an impending (non-Markov) switch to the worst equilibrium reduces the agent's future payoffs, and hence reduces the agent's current incentive cost. But the eventual switch appears to squander surplus, and it is less clear how this can make *both* parties better off. The cases in which both parties benefit from such front-loading are those in which the Markov equilibrium features some delay. Front-loading effectively pushes the agent's effort forward, coupling more intense initial effort with the eventual switch to the undesirable equilibrium, which can be sufficiently surplus-enhancing as to allow both parties to benefit.

Third, our analysis brings out the role of bargaining power in dynamic agency problems. A useful benchmark is provided by Bergemann and Hege [2], who examine an

analogous model in which the agent rather than the principal makes the offers in each period and hence has the bargaining power in the relationship.² The comparison is somewhat hindered by some technical difficulties overlooked in their analysis. As we explain in Appendix A.9, the same forces that sometimes preclude the existence of Markov equilibria in our model also arise in Bergemann and Hege [2]. Similarly, in other circumstances, multiple Markov equilibria may arise in their model (as in ours). As a result, they cannot impose their restriction to pure strategies, or equivalently, to beliefs of the principal and agent that coincide, and their results for the non-observable case are formally incorrect: their “Markov equilibrium” is not fully specified, does not always exist, and need not be unique. The characterization of the set of “Markov equilibria” in their model, as well as the investigation of non-Markov equilibria, remains an open question. However, we believe that the outcome they describe is indeed the outcome of a perfect Bayesian equilibrium satisfying a requirement similar to our recursive Markov requirement, and we conjecture that the set of such (recursive Markov) equilibrium outcomes converges to a unique limit as frequencies increase.³

Using this outcome as the point of comparison, there are some similarities in results across the two papers. Bergemann and Hege find four types of equilibrium behavior, each existing in a different region of parameter values. We similarly identify four analogous regions of parameter values (though not the same). But there are surprising differences: most notably, the agent might prefer that the bargaining power rest with the principal. The future offers made by an agent may be sufficiently lucrative that the agent can credibly claim to work in the current period only if those future offers are inefficiently delayed, to the extent that the agent fares better under the principal’s less generous but undelayed offers. However, one must bear in mind that these comparisons apply to “Markov” equilibria only, the focus of their analysis; the non-Markov equilibria that we examine exhibit properties quite different from those of recursive Markov equilibria.

The comparison is developed in detail in Section 4, which also provides additional benchmarks. We consider a model in which the principal can observe the agent’s effort, so there is no hidden information problem, identifying circumstances in which this observability makes the principal worse off. We also consider a variation of the model in which there is no learning. The model is then stationary, with each failure leading to a continuation game identical to the original game, but with non-Markov equilibria still giving rise to payoffs that are unattainable under Markov equilibria. We next consider

²Bergemann, Hege and Peng [3] present an alternative model of sequential investment in a venture capital project, without an agency problem, which they then use as a foundation for an empirical analysis of venture capital projects.

³Verifying these statements would involve arguments similar to those we have used, but this is not a direct translation. We suspect the corresponding limiting statement regarding uniqueness of perfect Bayesian equilibrium outcomes does not hold in their paper; as in our paper, one might be able to leverage the multiplicity of equilibria in the discrete-time game to construct non-Markov equilibrium outcomes that are distinct in the limit.

a model in which the principal has commitment power, identifying circumstances under which the ability to commit does and does not allow the principal to garner higher payoffs.

1.3 Applications

We view this model as potentially useful in examining a number of applications. The leading one is the case of a venture capitalist who must advance funds to an entrepreneur who is conducting experiments potentially capable of yielding a valuable innovation. A large empirical literature has studied venture capital. This literature makes it clear that actual venture capital contracts are more complicated than those captured by our model, but we also find that the structural features, methodological issues, and equilibrium characteristics of our model appear prominently in this literature.

The three basic structural elements of our model are moral hazard, learning, and the absence of commitment. The first two appear prominently in Hall’s [16] summary of the venture capital literature, that emphasizes the importance of (i) the venture capitalist’s inability to perfectly monitor the hidden actions of the agent, giving rise to moral hazard and (ii) learning over time about the potential of the project.⁴ Kaplan and Strömberg [20, Section 4.1, pp. 295–299] use Holmström’s principal-agent model (as do we) as the point of departure for assessing the financing provisions of venture capital contracts, arguing (p. 296) that their analysis is “largely consistent with both the assumptions and predictions of the classical principal-agent approach.” Kaplan and Strömberg [20, p. 313] also report that venture capital contracts are frequently renegotiated, reflecting the principal’s inability to commit.

Hall [16] identifies another basic feature of the venture capital market as rates of return for the venture capitalist that exceed those normally used for conventional investment. The latter feature, which distinguishes our analysis from Bergemann and Hege [2] (whose principal invariably earns a zero payoff, while our principal’s payoff may be positive), is well-documented in the empirical literature (see, for instance, Blass and Yosha [7]). This reflects the fact that funding for project development is scarce: technology managers often report that they have more projects they would like to undertake than funds to spend on them.⁵

⁴In the words of Hall [16, p. 411], “An important characteristic of uncertainty for the financing of investment in innovation is the fact that as investments are made over time, new information arrives which reduces or changes the uncertainty. The consequence of this fact is that the decision to invest in any particular project is not a once and for all decision, but has to be reassessed throughout the life of the project. In addition to making such investment a real option, the sequence of decisions complicates the analysis by introducing dynamic elements into the interaction of the financier (either within or without the firm) and the innovator.”

⁵See Peeters and van Pottelsberghe [24]; Jovanovic and Szentes [19] stress the scarcity of venture capital.

The methodological difficulties in our analysis arise to a large extent out of the possibility that the unobservability of the agent’s actions can cause the beliefs of the principal and agent to differ. Hall [16] points to this asymmetric information as another fundamental theme of the principal-agent literature, while Cornelli and Yosha [10, p. 4] note that, “At the time of initial venture capital financing, the entrepreneur and the financier are often equally informed regarding the project’s chances of success, and the true quality is gradually revealed to both. The main conflict of interest is the asymmetry of information about future actions of the entrepreneur.”

Finally, our finding that equilibrium effort is front-loaded resonates with the empirical literature. The common characteristic of our efficient equilibria is an initial (or front-loaded) sequence of funding and effort followed by premature (from the agent’s point of view) downsizing or termination of investment. It is well recognized in the theoretical and empirical literature that venture capital contacts typically provide funding in stages, with new funding following the arrival of new information (e.g., Admati and Pfleiderer [1], Cornelli and Yosha [10], Gompers [14], Gompers and Lerner [15] and Hall [16]). Moreover, Admati and Pfleiderer [1] characterize venture capital as a device to force the termination of projects entrepreneurs would otherwise prefer to continue, while Gompers [14, p.2] notes that venture capitalists maintain “the option to periodically abandon projects.”

Moreover, the empirical literature lends support to the mechanisms behind the equilibrium features in our paper. While one can conceive of alternative explanations for termination or downsizing, the role of the dynamic agency cost seems to be well-recognized in the literature: Cornelli and Yosha [10, p.1] explain that “the option to abandon is essential because an entrepreneur will almost never quit a failing project as long as others are providing capital” and find that investors often wish to downscale or terminate projects that entrepreneurs are anxious to continue. The role of termination or downsizing in mitigating this cost is also stressed by the literature: Sahlman [27, p. 507] states that “The credible threat to abandon a venture, even when the firm might be economically viable, is the key to the relationship between the entrepreneur and the venture capitalist....” Downsizing, which has generally attracted less attention than termination, is studied by Denis and Shome [11]. These features point to the (constrained efficient) non-Markov equilibrium as a more realistic description of actual relationships than the Markov equilibria usually considered in the theoretical literature. Our analysis establishes that such downsizing or termination requires no commitment power, lending support to Sahlman’s [27] statement above that such threats can be credible—though as we show, credibility considerations may ensure that downsizing (rather than termination) is the only option in the absence of commitment.

1.4 Related Literature

As we explain in more detail in Section 2, our model combines a standard hidden-action principal agent model (as introduced by Holmström [18]) with a standard bandit model of learning (as introduced in economics by Rothschild [26]). Each of these models in isolation is canonical and has been the subject of a large literature.⁶ Their interaction leads to interesting effects that are explored in this paper.

Other papers (in addition to Bergemann and Hege [2]) have also examined repeated-principal agent problems. Gerardi and Maestri [13] examine a model in which a principal must make a choice whose payoff depends on the realization of an unknown state, and who hires an agent to exert costly but unobservable effort in order to generate a signal that is informative about the state. In contrast to our model, the principal need not provide funding to the agent in order for the latter to exert effort, the length of the relationship is fixed, the outcome of the agent's experiments is unobservable, and the principal can ultimately observe and condition payments on the state.

Mason and Välimäki [23] examine a model in which the probability of a success is known and the principal need not advance the cost of experimentation to the agent. The agent has a convex cost of effort, creating an incentive to smooth effort over time. The principal makes a single payment to the agent, upon successful completion of the project. If the principal is unable to commit, then the problem and the agent's payment are stationary. If able to commit, the principal offers a payment schedule that declines over time in order to counteract the agent's effort-smoothing incentive to push effort into the future.

2 The Model

2.1 The Agency Relationship

2.1.1 Actions

We consider a long-term interaction between a principal (she) and an agent (he). The agent has access to a project that is either good or bad. The project's type is unknown, with principal and agent initially assigning probability $\bar{q} \in [0, 1)$ to the event that it is good. The case $\bar{q} = 1$ requires minor adjustments in the analysis, and is summarized in Section 4.2.

The game starts at time $t = 0$ and takes place in continuous time. At time t , the principal makes a (possibly history-dependent) offer s_t to the agent, where $s_t \in [0, 1]$

⁶See Bolton and Dewatripont [8] and Martimort and Laffont [22] for introductions to the principal-agent literature, and Berry and Fristedt [5] for a survey of bandit models of learning.

identifies the principal’s share of the proceeds from the project.⁷ Whenever the principal makes an offer to the agent, she cannot make another offer until $dt > 0$ units of time have passed. This inertia is essential to the rigorous analysis of the model, as explained in Section 1.2. We present a formal description of the game with inertia in Appendix A.2. What follows is a heuristic description, with emphasis on the limiting case as $dt \rightarrow 0$.

Whenever an offer is made, the principal advances the amount cdt to the agent, and the agent immediately decides whether to conduct an experiment, at cost cdt , or to shirk. If the experiment is conducted and the project is bad, the result is inevitably a failure, yielding no payoffs but leaving open the possibility of conducting further experiments. If the project is good, the experiment yields a success with probability pdt and a failure with probability $1 - pdt$, where $p > 0$. Alternatively, if the agent shirks, there is no success, and the agent expropriates the advance cdt .

The game ends at time t if and only if there is a success at that time. A success constitutes a breakthrough that generates a surplus of π , representing the future value of a successful project and obviating the need for further experimentation. The principal receives payoff πs_t from a success and the agent retains $\pi(1 - s_t)$. The principal and agent discount at the common rate r . There is no commitment on either side.

In the baseline model, the principal cannot observe the agent’s action, observing only a success (if the agent experiments and draws a favorable outcome) or failure (otherwise).⁸ We investigate the case of observable effort and the case of commitment by the principal in Sections 4.1 and 4.3.

We have described our principal-agent model as one in which the agent can divert cash that is being advanced to finance the project. In doing so, we follow Biais, Mariotti, Plantin and Rochet [6], among many others. However, we could equivalently describe the model as the simplest case—two actions and two outcomes—of the canonical hidden-action model (Holmström [18]). Think of the agent as either exerting low effort, which allows the agent to derive utility c from leisure, or high effort, which precludes such utility. There are two outcomes \bar{y} and $\underline{y} < \bar{y}$. Low effort always leads to outcome \underline{y} while high effort yields outcome \bar{y} with probability p . The principal observes only outcomes, and can attach payment $\bar{t} \in \mathbb{R}_+$ to outcome \bar{y} and payment $\underline{t} \in \mathbb{R}_+$ to outcome \underline{y} . The principal’s payoff is the outcome minus the transfer, while the agent’s payoff is the transfer plus the value of leisure. We then need only note that the optimal contract necessarily sets $\underline{t} = 0$, and then set $\bar{y} = \pi - c$, $\underline{y} = -c$, and $\bar{t} = \pi(1 - s)$ to render the models precisely equivalent.⁹

We place our principal-agent model in the simplest case of the canonical bandit model of learning—there are two arms, one of which is constant, and two signals, one of which is

⁷The restriction to the unit interval is without loss.

⁸This is the counterpart of Bergemann and Hege’s [2] “arm’s length” financing.

⁹It is then a simple change of reference point to view the agent as deriving no utility from low effort and incurring cost c of high effort.

fully revealing: a success cannot obtain when the project is bad. We could alternatively have worked with a “bad news” model, in which failures can only obtain when the project is bad. The mathematical difficulties in going beyond the good news or bad news bandit model when examining strategic interactions are well-known, prompting the overwhelming concentration in the literature on such models (though see Keller and Rady [21] for an exception).

2.1.2 Strategies and Equilibrium

Appendix A.2 provides the formal definitions of strategies and outcomes. Intuitively, the principal’s behavior strategy σ^P specifies at every instant, as a function of all the public information h_t^P (namely all the past offers), a choice of either an offer to make or to delay making an offer. Implicitly, we restrict attention to the case in which all experiments have been unsuccessful, as the game ends otherwise. The prospect that the principal might delay an offer is first encountered in Section 3.1.2, where we explain how this is a convenient device for modeling the prospect that the project might be downsized. A behavior strategy for the agent σ^A maps all information h_t , both public (past offers and the current offer to which the agent is responding) and private (past effort choices), into a decision to work or shirk, given the outstanding offer by the principal.

We examine weak perfect Bayesian equilibria of this game. In addition, because actions by the agent are not observed and the principal does not know the state, it is natural to impose the “no signaling what you don’t know” requirement on posterior beliefs after histories h_t (resp. h_t^P for the principal) that have probability zero under $\sigma = (\sigma^P, \sigma^A)$. In particular, consider a history h_t for the agent that arises with zero probability, either because the principal has made an out-of-equilibrium offer or the agent an out-of-equilibrium effort choice. There is some strategy profile (σ'^P, σ'^A) under which this history would be on the equilibrium path, and we assume the agent holds the belief that he would derive under Bayes’ rule given the strategy profile (σ'^P, σ'^A) . Indeed, there will be many such strategy profiles consistent with the agent’s history, but all of them feature the same sequence of effort choices on the part of the agent (namely, the effort choices the agent actually made), and so all give rise to the same belief. Similarly, given a history h_t^P for the principal that arises with zero probability (because the principal made an out-of-equilibrium offer), the principal holds the belief that would be derived from Bayes’ rule under the probability distribution induced by any strategy profile (σ'^P, σ^A) under which this history would be on the equilibrium path. Note that we hold σ^A fixed at the agent’s equilibrium strategy here, since the principal can observe nothing that is inconsistent with the agent’s equilibrium strategy. Again, any such (σ'^P, σ^A) specifies the same sequence of offers from the principal and the same response from the agent, and hence the same belief. “Equilibrium” henceforth refers to a weak perfect Bayesian Equilibrium satisfying these requirements. A class of equilibria of particular interest are recursive Markov equilibria,

discussed next.

2.2 What is Wrong with Markov Equilibrium?

Our game displays both incomplete information and hidden actions. The agent knows all actions taken. Therefore, his belief is simply a number $q^A \in [0, 1]$, namely the subjective probability that he attaches to the project being good. By definition of an equilibrium, this is a function of his history of experiments only. It is payoff-relevant, as it affects the expected value of following different courses of action. Therefore, any reasonable definition of Markov equilibrium must include q^A as a state variable. This posterior belief is continually revised throughout the course of play, with each failure being bad news, leading to a more pessimistic posterior expectation that the project is good.

The agent's hidden action gives rise to hidden information: if the agent deviates, he will update his belief unbeknownst to the principal, and this will affect his future incentives to work, given the future equilibrium offers, and hence his payoff from deviating. In turn, the principal must compute this payoff in order to determine which offers will induce the agent to work. Therefore, the principal's belief *about the agent's belief*, $q^P \in \Delta[0, 1]$ is payoff-relevant. To be clear, q^P is a distribution over the agent's belief: if the agent randomizes his effort decision, the principal will have uncertainty regarding the agent's posterior belief, as she does not know the realized effort choice. Fortunately, her belief q^P is commonly known.

The natural state variable for a Markov equilibrium is thus the pair (q^A, q^P) . If the agent's equilibrium strategy were pure, then as long as the agent does not deviate, those two beliefs would coincide, and we could use the single belief q^A as a state variable. This is the approach followed by Bergemann and Hege [2]. Unfortunately, two difficulties arise. First, when the agent evaluates the benefit of deviating, he must consider how the game will evolve after such a deviation, from which point the beliefs q^A and q^P will differ. More importantly, there is typically no equilibrium in which the agent's strategy is pure. Here is why.

Expectations about the agent's action affect his continuation payoff, as they affect the principal's posterior belief. Hence, we can define a share s^S for which the agent would be indifferent between working and shirking if he is expected to shirk; and a share s^W for which he is indifferent if he is expected to work. If there was no learning ($\bar{q} = 1$), these shares would coincide, as the agent's action would not affect the principal's posterior belief and hence continuation payoffs. But if $\bar{q} < 1$ they do not coincide. This gives rise to two possibilities that both can arise depending on the parameters. If $s^S < s^W$, there are multiple equilibria: for each share \bar{s} in the interval $[s^S, s^W]$, we can construct an equilibrium in which the agent shirks if $s > \bar{s}$, and works if $s \leq \bar{s}$. Given that $s^S < s^W$, these beliefs are consistent with the agent's incentives. On the other hand, if $s^S > s^W$, then there is no pure best-reply for the agent if a share s in the interval (s^S, s^W) is

offered!¹⁰ If it were optimal to shirk and the principal would expect it, then the agent would not find it optimal to shirk after all (because $s < s^S$). Similarly, if it were optimal to work and the principal would expect it, then working could not be optimal after all, because $s > s^W$. As a result, whenever these thresholds are ordered in this way, as does occur for a wide range of parameters, then the agent must randomize between shirking and working.

Once the agent randomizes, q^P is no longer equal to q^A , as the principal is uncertain about the agent’s realized action. The analysis thus cannot ignore states for which q^P is non-degenerate (and so differs from q^A).

Unfortunately, working with the larger state (q^A, q^P) is still not enough. Markov equilibria—even on this larger state space—fail to exist. The reason is closely related to the non-existence of pure-strategy equilibria. If a share $s \in (s^S, s^W)$ is offered, the agent must randomize over effort choices, which means that he must be indifferent between two continuation games, one following work and one following shirk. This indifference requires strategies in these continuations to be “fine-tuned,” and this fine-tuning depends on the exact share s that was offered; this share, however, is not encoded in the later values of the pair (q^A, q^P) .

This is a common feature of extensive-form games of incomplete information (see for instance Fudenberg, Levine and Tirole [12] and Hellwig [17]). To restore existence while straying as little as possible from Markov equilibrium, Fudenberg, Levine and Tirole [12] introduce the concept of weak Markov equilibrium, afterwards used extensively in the literature. This concept allows behavior to depend not only on the current belief, but also on the preceding offer. Unfortunately, this does not suffice here, for reasons that are subtle. The fine-tuning that is mentioned in the previous paragraph cannot be achieved right after the offer s is made. As it turns out, it requires strategies to exhibit arbitrarily long memory of the exact offer that prompted the initial mixing.

This issue is discussed in detail in Appendix A, where we also define a solution concept, recursive Markov equilibrium, that is analogous to weak Markov equilibrium and is appropriate for our game. This concept coincides with Markov equilibrium (with generic state (q^A, q^P)) whenever it exists, or with weak Markov equilibrium whenever the latter exists. Roughly speaking, a recursive Markov equilibrium requires that (i) play coincides with a Markov equilibrium whenever such an equilibrium exists (which it does for low enough beliefs, as we prove), and (ii) if the state does not change from one period to the next, then neither do equilibrium actions.

As we show, recursive Markov equilibria always exist in our game. Moreover, as the minimum length of time between two consecutive offers vanishes (i.e., as $dt \rightarrow 0$), all weak Markov equilibria yield the same outcome, and the belief of the principal converges to the belief of the agent, so that, in the limit, we can describe strategies as if there was a

¹⁰Such offers do not appear along the equilibrium path, but nonetheless checking the optimality of an offer outside the interval requires knowing the payoff that would result from an offer within it.

single state variable $q = q^A = q^P$. Hence, considering this limit yields sharp predictions, as well as particularly tractable solutions.

As shown in Appendix A.9, cases in which $s^W < s^S$ arise and cases in which $s^W > s^S$ can also arise when the agent makes the offers instead of the principal, as in Bergemann and Hege [2]. Hence, Markov equilibria need not exist in their model; and when they do, they are not unique. Nevertheless, the patterns that emerge from their computations bear strong similarities with ours, and we suspect that, if one were to (i) study the recursive Markov equilibria of their game, and (ii) take the continuous-time limit of the set of equilibrium outcomes, one would obtain qualitatively similar results. Given these similarities, we view their results as strongly suggestive of what is likely to come out of such an analysis, and will therefore use their predictions to speculate about how bargaining power affects equilibrium play.

2.3 The First-Best Policy

Suppose there is no agency problem—either the principal can conduct the experiments (or equivalently the agent can fund the experiments), or there is no monitoring problem and hence the agent necessarily experiments whenever asked to do so.

The principal will experiment until either achieving a success, or being rendered sufficiently pessimistic by a string of failures as to deem further experimentation unprofitable. The optimal policy, then, is to choose an optimal stopping time, given the initial belief. That is, the principal chooses $T \geq 0$ so as to maximize the normalized expected value of the project, given by

$$V(\bar{q}) = \mathbb{E} \left[\pi e^{-rv} \mathbf{1}_{v \leq T} - \int_0^{v \wedge T} c e^{-rt} dt \right],$$

where r is the discount rate, v is the random time at which a success occurs, and $\mathbf{1}_E$ is the indicator of the event E .

The probability that no success has obtained by time t is $\exp\left(-\int_0^t p q_v dv\right)$, and we can then use the law of iterated expectations to rewrite this payoff as

$$V(\bar{q}) = \int_0^T e^{-rt - \int_0^t p q_v dv} (p q_t \pi - c) dt.$$

From this formula, it is clear that it is optimal to pick $T \geq t$ if and only if $p q_t \pi - c > 0$. Hence, the principal experiments if and only if

$$q_t > \frac{c}{p\pi}. \tag{1}$$

The optimal stopping time T then solves $q_T = c/p\pi$.

Appendix B develops an expression for the optimal stopping time that immediately yields some intuitive comparative statics. The first-best policy operates the project longer when the prior probability \bar{q} is larger (because it then takes longer to become so pessimistic as to terminate), when (holding p fixed) the benefit-cost ratio $p\pi/c$ is larger (because more pessimism is then required before abandoning the project), and when (holding $p\pi/c$ fixed) the success probability p is smaller (because consistent failure is then less informative).

3 Characterization of Equilibria

This section describes the set of equilibrium outcomes, characterizing both behavior and payoffs, in the limit as $dt \rightarrow 0$. This is the limit identified in Section 2.1. We explain the intuition behind these outcomes. Our intention is that this description, including some heuristic derivations, should be sufficiently compelling that most readers need not delve into the technical details behind this description. The formal arguments supporting this section’s results require a characterization of equilibrium behavior and payoffs for the game with inertia and a demonstration that the behavior and payoffs described here are the unique limits (as inertia dt vanishes) of such equilibria. This is the sense in which we use “uniqueness” in what follows. See Appendix A for details.

3.1 Recursive Markov Equilibria

3.1.1 No Delay

We begin by examining equilibria in which the principal never delays making an offer, or equivalently in which the project is never downsized, and the agent is always indifferent between accepting and rejecting the equilibrium offer.

Let q_t be the common (on path) belief at time t . This belief will be continually revised downward (in the absence of a game-ending success), until the expected value of continued experimentation hits zero. At the last instant, the interaction is a static principal-agent problem. The agent can earn cdt from this final encounter by shirking, and can earn $(1 - s)pq\pi dt$ by working. The principal will set the share s so that the agent’s incentive constraint binds, or $cdt = (1 - s)pq\pi dt$. Using this relationship, the principal’s payoff in this final encounter will then satisfy

$$(spq\pi - c)dt = (pq\pi - 2c)dt,$$

and the principal will accordingly abandon experimentation at the value \underline{q} at which this expression is zero, or

$$\underline{q} := \frac{2c}{p\pi}.$$

This failure boundary highlights the cost of agency. The first-best policy derived in Section 2.3 experiments until the posterior drops to $c/p\pi$, while the agency cost forces experimentation to cease at $2c/p\pi$. In the absence of agency, experimentation continues until the expected surplus $pq\pi$ just suffices to cover the experimentation cost c . In the presence of agency, the principal must not only pay the cost of the experiment c , but must also provide the agent with a rent of at least c , to ensure the agent does not shirk and appropriate the experimental funding. This effectively doubles the cost of experimentation, in the process doubling the termination boundary.

Now consider behavior away from the failure boundary. The belief q_t is the state variable in a recursive Markov equilibrium, and equilibrium behavior will depend only on this belief (because we are working in the limit as $dt \rightarrow 0$). Let $v(q)$ and $w(q)$ denote the “ex post” equilibrium payoffs of the principal and the agent, respectively, given that the current belief is q and that the principal has not yet made an offer to the agent. By ex post, we refer to the payoffs after the requisite waiting time dt has passed and the principal is on the verge of making the next offer. Let $s(q)$ denote the offer made by the principal at belief q , leading to a payoff $\pi s(q)$ for the principal and $\pi(1 - s(q))$ for the agent if the project is successful.

The principal’s payoff $v(q)$ follows a differential equation. To interpret this equation, let us first write the corresponding difference equation for a given $dt > 0$ (up to second order terms):

$$v(q_t) = (pq_t\pi s(q_t) - c)dt + (1 - rdt)(1 - pq_tdt)v(q_{t+dt}).$$

The first term on the right is the expected payoff in the current period, consisting of the probability of a success pq_t multiplied by the payoff $\pi s(q_t)$ in the event of a success, minus the cost of the advance c , all scaled by the period length dt . The second term is the continuation value to the principal in the next period $v(q_{t+dt})$, evaluated at next period’s belief q_{t+dt} and multiplied by the discount factor $1 - rdt$ and the probability $1 - pq_tdt$ of reaching the next period via a failure.

Taking the limit $dt \rightarrow 0$, we get the differential equation corresponding to the frictionless limit,

$$(r + pq)v(q) = pq\pi s(q) - c - pq(1 - q)v'(q). \quad (2)$$

The left side is the annuity on the project, given the effective discount factor $(r + pq)$. This must equal the sum of the flow payoff, $pq\pi s(q) - c$, and the capital loss, $v'(q)\dot{q}$, imposed by the deterioration of the posterior belief induced by a failure. To see how \dot{q} appears in (2), we note that Bayes’ rule gives

$$q_{t+dt} = \frac{q_t(1 - pdt)}{1 - pq_tdt}, \quad (3)$$

from which it follows that, in the limit,

$$\dot{q}_t = -pq_t(1 - q_t). \quad (4)$$

Similarly, the payoff to the agent, $w(q_t)$, must solve, to the second order,

$$\begin{aligned} w(q_t) &= pq_t\pi(1 - s(q_t))dt + (1 - rdt)(1 - pq_tdt)w(q_{t+dt}) \\ &= cdt + (1 - rdt)(w(q_{t+dt}) + x(q_t)dt). \end{aligned} \quad (5)$$

The first equality gives the agent's equilibrium value as the sum of the agent's current-period payoff $pq_t\pi(1 - s(q_t))dt$ and the agent's continuation payoff $w(q_{t+dt})$, discounted and weighted by the probability the game does not end. The second equality is the agent's incentive constraint. The agent must find the equilibrium payoff at least as attractive as the alternative of shirking. The payoff from shirking includes the appropriation of the experimentation cost cdt , plus the discounted continuation payoff $w(q_{t+dt})$, which is now received with certainty and is augmented by $x(q_t)$, defined to be the marginal gain from $t + dt$ onward from shirking at time t unbeknownst to the principal. The agent's incentive constraint must bind in a recursive Markov equilibrium, since otherwise the principal could increase her share without affecting the agent's effort.

To evaluate $x(q_t)$, note that, when shirking, the agent holds an unchanged posterior belief, q_t , while the principal wrongly updates to $q_{t+dt} < q_t$. If the equilibrium expectation is that the agent works in all subsequent periods, then he will do so as well if he is more optimistic. Furthermore, the agent's value (when he always works) arises out of the induced probability of a success in the subsequent periods. A success in a subsequent period occurs with a probability that is proportional to his current belief. As a result, the value from being more optimistic is q_t/q_{t+dt} higher than if he had not deviated. Hence,

$$x(q_t)dt = \left(\frac{q_t}{q_{t+dt}} - 1 \right) w(q_{t+dt}), \quad (6)$$

or, taking the limit $dt \rightarrow 0$ and using (4),

$$x(q_t) = p(1 - q_t)w(q_t).$$

Using this expression and again taking the limit $dt \rightarrow 0$, the agent's payoff satisfies

$$\begin{aligned} 0 &= pq\pi(1 - s(q)) - pq(1 - q)w'(q) - (r + pq)w(q) \\ &= c - pq(1 - q)w'(q) - (r + pq)w(q) + pw(q). \end{aligned} \quad (7)$$

The term $pw(q)$ in (7) reflects the future benefit from shirking now. This gives rise to what we call a *dynamic agency cost*. One virtue of shirking is that it ensures the game continues, rather than risking a game-ending success. The larger the agent's continuation value $w(q)$, the larger the temptation to shirk, and hence the more expensive will the principal find it to induce effort.

This gives us three differential equations (equation (2) and the two equalities in (7)) in three unknown functions (v , w and s). We have already identified the relevant boundary

condition, namely that experimentation ceases when the posterior q drops to \underline{q} . We can accordingly solve for the candidate equilibrium.

It is helpful to express this solution in terms of the following notation:

$$\psi := \frac{p\pi - 2c}{c} \quad \sigma := \frac{p}{r}.$$

Noting that $\psi = p\pi/c - 2$, we can interpret ψ as the benefit-cost ratio of the experiment, or equivalently as a measure of the size of the potential surplus. We accordingly interpret high values of ψ as identifying high-surplus projects and low values of ψ as identifying low-surplus projects. We assume that ψ is positive, since otherwise $\underline{q} = 2/(2 + \psi)$ is greater than 1, and hence no experimentation takes place no matter what the prior belief. Notice next that σ is larger as r is smaller, and hence the players are more patient. We thus refer to large values of σ as identifying “patient projects” and low values of σ as identifying impatient projects.

Using this notation, we can use the second equation from (7) to solve for w , using as boundary condition $w(\underline{q}) = 0$:

$$w(q) = \frac{\frac{q}{\underline{q}}(1 - q\sigma) \left(\frac{(1-q)\underline{q}}{q(1-\underline{q})} \right)^{\frac{1}{\sigma}} - (1 - q\sigma) \frac{c}{r}}{\sigma - 1}. \quad (8)$$

Figure 2 (below) illustrates this solution.

It is a natural expectation that $w(q)$ should increase in q , since the agent seemingly always has the option of shirking and the payoff from doing so increases with the time until experimentation stops. Figure 2 shows that $w(q)$ may decrease in q for large values of q . To see how this might happen, fix a positive period-of-inertia length dt , so that the principal will make a finite number of offers before terminating experimentation, and consider what happens to the agent’s value if the prior probability \bar{q} is increased just enough to ensure that the maximum number of experiments has increased by one. From the incentive constraint (5) and (6) we see that this extra experimentation opportunity (i) gives the agent a chance to expropriate the cost of experimentation c (which the agent will not do in equilibrium, but nonetheless is indifferent between doing so and not), (ii) delays the agent’s current value by one period and hence discounts it, and (iii) increases this current value by a factor of the form q_t/q_{t+dt} , reflecting the agent’s more optimistic prior. The first and third of these are benefits, the second is a cost. The benefits will often outweigh the costs, for all priors, and W will then be increasing in q . However, the factor q_t/q_{t+dt} is smallest for large q , and hence if w is ever to be decreasing, it will be so for large q , as in Figure 2.

We can use the first equation from (7) to solve for $s(q)$, and then solve (2) for the

value to the principal, which gives, given the boundary condition $v(\underline{q}) = 0$,

$$v(q) = \left[\left(\frac{(1-q)\underline{q}}{q(1-\underline{q})} \right)^{\frac{1}{\sigma}} \left(1 - \frac{(1-q)\underline{q}(\psi+1)}{(1-\underline{q})(\sigma+1)} + \frac{q(1-\underline{q})}{\underline{q}(\sigma-1)} \right) + \frac{2(1-q\sigma)}{\sigma^2-1} - \frac{\sigma-q(\psi+2)}{\sigma+1} \right] \frac{c}{r}. \quad (9)$$

This seemingly complicated expression is actually straightforward to manipulate. For instance, a simple Taylor expansion reveals that v is approximately proportional to $(q-\underline{q})^2$ in the neighborhood of \underline{q} , while w is approximately proportional to $(q-\underline{q})$. Both parties' payoffs thus tend to zero as q approaches \underline{q} , since the net surplus $pq\pi - 2c$ declines to zero. The principal's payoff tends to zero faster, as there are two forces behind this disappearing payoff: the remaining time until experimentation stops for good vanishes, and the mark-up she gets from success does so as well. The agent's mark-up, on the other hand, does not disappear, as shirking yields a benefit that is independent of q , and hence the agent's payoff is proportional to the remaining amount of experimentation time.

These strategies constitute an equilibrium if and only if the principal's participation constraint $v(q) \geq 0$ is satisfied for $q \in [\underline{q}, \bar{q}]$. (The agent's incentive constraint implies the corresponding participation constraint for the agent.) First, for $q = 1$, (9) immediately gives:

$$v(1) = \frac{\psi - \sigma c}{\sigma + 1 r}, \quad (10)$$

which is positive if and only if $\psi > \sigma$. This is the first indication that our candidate no-delay strategies will not always constitute an equilibrium.

To interpret the condition $\psi > \sigma$, let us rewrite (10) as

$$v(1) = \frac{\psi - \sigma c}{\sigma + 1 r} = \frac{p\pi - c}{r + p} - \frac{c}{r}. \quad (11)$$

When $q = \bar{q} = 1$, the project is known to be good, and there is no learning. Our candidate strategies will then operate the project as long as it takes to obtain a success. The first term on the right in (11) is the value of the surplus, calculated by dividing the (potentially perpetually received) flow value $p\pi - c$ by the effective discount rate of $r + p$, with r capturing the discounting and p capturing the hazard of a flow-ending success. The second term in (11) is the agent's equilibrium payoff. Since the agent can always shirk, ensuring that the project literally endures forever, the agent's payoff is the flow value c of expropriating the experimental advance divided by the discount rate r .

As the players become more patient (r decreases), the agent's equilibrium payoff increases without bound, as the discounted value of the payoff stream c becomes arbitrarily valuable. In contrast, the presence of p in the effective discount rate $r + p$, capturing the event that a success ends the game, ensures that the value of the surplus cannot similarly increase without bound, no matter how patient the players. But then the principal's payoff (given $\bar{q} = 1$), given by the difference between the value of the surplus and the agent's

payoff, can be positive only if the players are not too patient. The players are sufficiently impatient that $v(1) > 0$ when $\psi > \sigma$, and too patient for $v(1) > 0$ when $\psi < \sigma$. We say that we are dealing with impatient players (or an impatient project or simply impatience), in the former case, and patient players in the latter case.

We next examine the principal's payoff near \underline{q} . We have noted that $v(\underline{q}) = v'(\underline{q}) = 0$, so everything here hinges on the second derivative $v''(\underline{q})$. We can use the agent's incentive constraint (7) to eliminate the share s from (2) and then solve for

$$v' = \frac{pq\pi - 2c - pw - (r + pq)v}{pq(1 - q)}.$$

With some fortuitous foresight, we first investigate the derivative $v''(\underline{q})$ for the case in which $\psi = 2$. This case is particularly simple, as $\psi = 2$ implies $\underline{q} = 1/2$, and hence $pq(1 - q)$ is maximized at \underline{q} . Marginal variations in q will thus have no effect on $pq(1 - q)$, and we can take this product to be a constant. Using $v'(\underline{q}) = 0$ and calculating that $w'(\underline{q}) = c/(pq(1 - \underline{q}))$, we have

$$v''(\underline{q}) = p\pi - pw'(\underline{q}) = p\pi - p\frac{c}{pq(1 - \underline{q})}.$$

Hence, as q increases above \underline{q} , v' tends to increase in response to the increased value of the surplus (captured by $p\pi$), but to decrease in response to the agent's larger payoff ($-pw'(\underline{q})$). To see which force dominates, multiply by $pq(1 - \underline{q})$ and then use the definition of ψ to obtain

$$pq(1 - \underline{q})v''(\underline{q}) = pq\psi c - pc = pc\left(\frac{\psi}{2} - 1\right) = 0.$$

Hence, at $\psi = 2$, the surplus-increasing and agent-payoff-increasing effects of an increase in q precisely balance, and $v''(\underline{q}) = 0$. It is intuitive that larger values of ψ enhance the surplus effect, and hence $v''(\underline{q}) > 0$ for $\psi > 2$.¹¹ In this case, $v(\underline{q}) > 0$ for values of q near \underline{q} . We refer to these as high-surplus projects. Alternatively, smaller values of ψ attenuate the surplus effect, and hence $v''(\underline{q}) < 0$ for $\psi < 2$. In this case, $v(\underline{q}) < 0$ for values of q near \underline{q} . We refer to these as low-surplus projects.

This gives us information about the endpoints of the interval $[\underline{q}, 1]$ of possible posteriors. It is a straightforward calculation that v admits at most one inflection point, so that it is positive everywhere if it is positive at 1 and increasing at $q = \underline{q}$. We can then summarize our results as follows, with Lemma 9 in Appendix A.6.3 providing the corresponding formal argument:

- v is positive for values of $q > \underline{q}$ close to \underline{q} if $\psi > 2$, and negative if $\psi < 2$.

¹¹Once ψ edges over 2, we can no longer take $pq(1 - q)$ to be approximately constant in q , yielding a more complicated expression for v'' that confirms this intuition. In particular, $v''(\underline{q}) = \frac{(\psi+2)^3(\psi-2)c}{4\sigma\psi^2 r}$.

- $v(1)$ is positive if $\psi > \sigma$ and negative if $\psi < \sigma$.
- Hence, if $\psi > 2$ and $\psi > \sigma$, then $v(q) \geq 0$ for all $q \in [\underline{q}, 1]$.

The key implication is that our candidate equilibrium, in which the project is never downsized, is indeed an equilibrium only in the case of a high-surplus ($\psi > 2$), impatient ($\psi > \sigma$) project. In all other cases, the principal's payoff falls below zero for some beliefs.

3.1.2 Delay

If either $\psi < 2$ or $\psi < \sigma$, then the strategy profiles developed in the preceding subsection cannot constitute an equilibrium, as they yield a negative principal payoff for some beliefs. Whether it occurs at high or low beliefs, this negative payoff reflects the dynamic agency cost. The agent's continuation value is sufficiently lucrative, and hence shirking in order to ensure the continuation value is realized is sufficiently attractive, that the principal can induce the agent to work only at such expense as to render the principal's payoff negative. Equilibrium then requires that the principal make continuing the project less attractive for the agent, making it less attractive to shirk and hence reducing the cost to the principal of providing incentives.

We think of the principal as downsizing the project in order to reduce its attractiveness. In the underlying game with inertial periods of length dt between offers, we capture this downsizing by having the principal delay by more than the required time dt between offers. This slows the rate of experimentation per unit of time, and in the frictionless limit appears as if the project is operated on a smaller scale.

One way of capturing this behavior in a stationary limiting strategy is to require the principal to undertake a series of independent randomizations between making an offer or not. Conditional on not making an offer at time t , the principal's belief does not change, so that she randomizes at time $t+dt$ in the same way as at time t . The result is a reduced rate of experimentation per unit of time that effectively operates the project on a smaller scale. However, as is well known, such independent randomizations cannot be formally defined in the limit $dt \rightarrow 0$. Instead, there exists a payoff-equivalent way to describe them that can be formally formulated. To motivate it, suppose that the principal makes a constant offer s with probability β over each interval of length dt . This means that the total discounted present value of payoffs is proportional to $\beta s/r$, as randomization effectively scales the payoff flows by a factor of β . This not only directly captures the idea that the project is downsized by factor β , but leads nicely to the next observation, namely that this is exactly the same value that would be obtained absent any randomization if the discount rate were set to $r/\beta \geq r$ rather than r . Accordingly, we will not have the principal randomize, but control the rate of arrival of offers, or equivalently the discount rate, which is well-defined and payoff-equivalent. Letting $\lambda = 1/\beta$, we replace the discount factor r with an effective discount factor $r\lambda(q)$, where λ is controlled by the principal.

Because $\beta \leq 1$, we have $\lambda(q) \geq 1$, with $\lambda(q) = 1$ whenever there is no delay (as is the case throughout Section 3.1.1), and $\lambda(q) > 1$ indicating delay. The principal can obviously choose different amounts of delay for different posteriors, making λ a function of q , but the lack of commitment only allows her to choose $\lambda > 1$ when she is indifferent between delaying or not. Given that we have built the representation of this behavior into the discount factor, we will refer to the behavior as delay, remembering that it has an equivalent interpretation as downsizing.¹²

We must now rework the system of differential equations from Section 3.1.1 to incorporate delay. It can be optimal for the principal to delay only if the principal is indifferent between receiving the resulting payoff later rather than sooner. This in turn will be the case only if the principal's payoff is identically zero, so whenever there is delay we have $v = v' = 0$. In turn, the principal's payoff is zero at q_t and at q_{t+dt} only if her flow payoff at q_t is zero, which implies

$$pq s(q) \pi = c, \quad (12)$$

and hence fixes the share $s(q)$. To reformulate equation (7), identifying the agent's payoff, let $w(q_t)$ identify the agent's payoff at posterior q_t . We are again working with ex post valuations, so that $w(q_t)$ is the agent's value when the principal is about to make an offer, given posterior q_t , and given that the inertial period dt as well as any extra delay has occurred. The discount rate r must then be replaced by $r\lambda(q)$. Combining the second equality in (7) with (12), we have

$$w(q) = \frac{pq\pi - 2c}{p}. \quad (13)$$

This gives $w'(q) = \pi$ which we can insert into the first equality of (7) (replacing r with $r\lambda(q)$) to obtain

$$(r\lambda(q) + pq)w(q) = pq^2\pi - c.$$

We can then solve for the delay

$$\lambda(q) = \frac{(2q - 1)\sigma}{q(\psi + 2) - 2}, \quad (14)$$

which is strictly larger than one if and only if

$$q(2\sigma - \psi - 2) > \sigma - 2. \quad (15)$$

¹²Bergemann and Hege [2] allow the principal to directly choose the scale of funding for an experiment from an interval $[0, \bar{c}]$ in each period, rather than simply choosing to fund the project or not. Lower levels of funding give rise to lower probabilities p of success. Lower values of c then have a natural interpretation as downsizing the experimentation. This convention gives rise to more complicated belief updating that becomes intractable when specifying out-of-equilibrium behavior.

We have thus solved for the values of both players' payoffs (given by $v(q) = 0$ and (13)), and for the delay over any interval of time in which there is delay (given by (14)).

From (15), note that the delay λ strictly exceeds 1 at $q = 1$ if and only if $\psi < \sigma$ and at $\underline{q} = 2/(2 + \psi)$ if and only if $\psi < 2$. In fact, since the left side is linear in q , we have $\lambda(q) \geq 1$ for all $q \in [\underline{q}, 1]$ if $\psi < \sigma$ and $\psi < 2$. Conversely, there can be no delay if $\psi > \sigma$ and $\psi > 2$. This fits the conditions derived in Section 3.1.1.

3.1.3 Recursive Markov Equilibrium Outcomes: Summary

We now have descriptions of two regimes of behavior, one without delay and one with delay. We must patch these together to construct equilibria. If $\psi > \sigma$ and $\psi > 2$, then from (15) there is no delay at $q = 1$ and no delay at $\underline{q} = 2/(2 + \psi)$. Further, since the left side of (15) is linear in q , we have no delay for any $q \in [\underline{q}, 1]$, matching the no-delay conditions derived in Section 3.1.1.

If $\psi < 2$ (but $\psi > \sigma$) it is natural to expect delay for low beliefs, with this delay disappearing as we reach the point at which equation (14) exactly gives no delay. That is, delay should disappear for beliefs above

$$q^{**} := \frac{2 - \sigma}{2 + \psi - 2\sigma}.$$

Alternatively, if $\psi < \sigma$ (but $\psi > 2$), we should expect delay to appear once the belief is sufficiently high for the function v defined by (2), which is positive for low q , to hit 0 (which it must, under these conditions). Because v has a unique inflection point, there is a unique value $q^* \in (\underline{q}, 1)$ that solves $v(q^*) = 0$.

We can summarize this with (with details in Appendix A.8):

Proposition 1 *Depending on the parameters of the problem, we have four types of recursive Markov equilibria, distinguished by their use of delay, summarized by:*

	<i>High Surplus</i> $\psi > 2$	<i>Low Surplus</i> $\psi < 2$
<i>Impatience, $\psi > \sigma$</i>	<i>No delay</i>	<i>Delay for low beliefs</i> $(q < q^{**})$
<i>Patience, $\psi < \sigma$</i>	<i>Delay for high beliefs</i> $(q > q^*)$	<i>Delay for all beliefs</i>

We can provide a more detailed description of these equilibria, with Appendix A.8 providing the technical arguments:

High Surplus, Impatience ($\psi > 2$ and $\psi > \sigma$): No Delay. In this case, there is no delay until the belief reaches \underline{q} , in case of repeated failures. At this stage, the project is abandoned. The relatively impatient agent does not value his future rents too highly, which makes it relatively inexpensive to induce him to work. Since the project produces a relatively high surplus, the principal's payoff from doing so is positive throughout. Formally, this is the case in which w and v are given by (8) and (9) and are both positive over the entire interval $[\underline{q}, 1]$.

High Surplus, Patience ($\psi > 2$ and $\psi < \sigma$): Delay for High Beliefs. In this case, the recursive Markov equilibrium is characterized by some belief $q^* \in (\underline{q}, 1)$. For higher beliefs, there is delay and the principal's payoff is zero. As the belief reaches q^* , delay disappears (taking a discontinuous drop in the process), and no further delay occurs until the project is abandoned (in the absence of an intervening success) when the belief reaches \underline{q} .

When beliefs are high, the agent expects a long-lasting relationship, which his patience renders quite lucrative, and effort is accordingly prohibitively expensive. Equilibrium requires delay in order to reduce the agent's continuation payoff and hence current cost. As the posterior approaches \underline{q} , the likely length of the agent's future rent stream declines, as does its value and hence the agent's current incentive cost. This eventually brings the relationship to a point where the principal can secure a positive payoff without delay.

Low Surplus, Impatience ($\psi < 2$ and $\psi > \sigma$): Delay for Low Beliefs. When beliefs are higher than q^{**} , there is no delay. When the belief reaches q^{**} , delay appears (with delay being continuous at q^{**}).

To understand why the dynamics are reversed, compared to the previous case, note that it is now not too costly to induce the agent to work when beliefs are high, since the impatient agent discounts the future heavily and does not anticipate a lucrative continuation payoff, and the principal here has no need to delay. However, when the principal becomes sufficiently pessimistic (q becomes sufficiently low), the low surplus generated by the project still makes it too costly to induce effort. The principal must then resort to delay in order to reduce the agent's cost and render her payoff nonnegative.

Low Surplus, Patience ($\psi < 2$ and $\psi < \sigma$): Perpetual Delay. In this case, the recursive Markov equilibrium involves delay for all values of $q \in [\underline{q}, 1]$. The agent's patience makes him relatively costly, and the low surplus generated by the project makes it relatively unprofitable, so that there is no belief at which the principal can generate a nonnegative payoff without delay. Formally, λ , as given by (14), is larger than one over $[\underline{q}, 1]$.

3.1.4 Recursive Markov Equilibrium Outcomes: Lessons

What have we learned from studying recursive Markov equilibria? First, the dynamic agency cost is a formidable force. A single encounter between the principal and agent would be profitable for the principal whenever $q > \underline{q}$. The dynamic agency cost increases the incentive cost of the agent to such an extent that only in the event of a high surplus, impatient project can the principal secure a positive payoff, no matter what the posterior belief, without resorting to cost-reducing delay.

Second, and consequently, a project that is more likely to be good is not necessarily better for the principal. This is obviously the case for a high surplus, patient project, where the principal is doomed to a zero payoff for high beliefs but earns a positive payoff when less optimistic. In this case, the lucrative future payoffs anticipated by an optimistic agent make it particularly expensive for the principal to induce current effort, while a more pessimistic agent anticipates a less attractive future (even though still patient), and can be induced to work at a sufficiently lower cost as to allow the principal a positive payoff. Moreover, even when the principal's payoff is positive, it need not be increasing in the probability the project is good. Figure 1 illustrates two cases (both high surplus projects) where this does not happen. A higher probability that the project is good gives rise to a higher surplus, but also makes creating incentives for the agent to work more expensive (because the agent anticipates a more lucrative future). The latter effect may overwhelm the former, and hence the principal may thus prefer to be pessimistic about the project. Alternatively, the principal may find a project with lower surplus more attractive than a higher-surplus project, especially (but not only) if the former is coupled with a less patient agent. The larger surplus may increase the cost of inducing the agent to work so much as to reduce the principal's expected payoff.

Third, low surplus projects lead to delay for low posterior beliefs. Do we reach the terminal belief \underline{q} in finite time, or does delay increase sufficiently fast, in terms of discounting, that the event that the belief reaches \underline{q} is essentially discounted into irrelevance? If the latter is the case, we would never observe a project being abandoned, but would instead see them wither away, limping along at continuing but negligible funding.

It is straightforward to verify that for low surplus projects, not only does $\lambda(q)$ diverge as $q \downarrow \underline{q}$, but so does the integral

$$\lim_{q \downarrow \underline{q}} \int_q^{\bar{q}} \lambda(v) dv.$$

That is, the event that the project is abandoned is entirely discounted away in those cases in which there is delay for low beliefs. This means that, in real time, the belief \underline{q} is only reached asymptotically, so that the project is never really abandoned. Rather, the pace of experimentation slows sufficiently fast that this belief is never reached. An analogous feature appears in models of strategic experimentation (see, for example, Bolton and Harris [9, p. 363] and Keller and Rady [21, p. 290]).

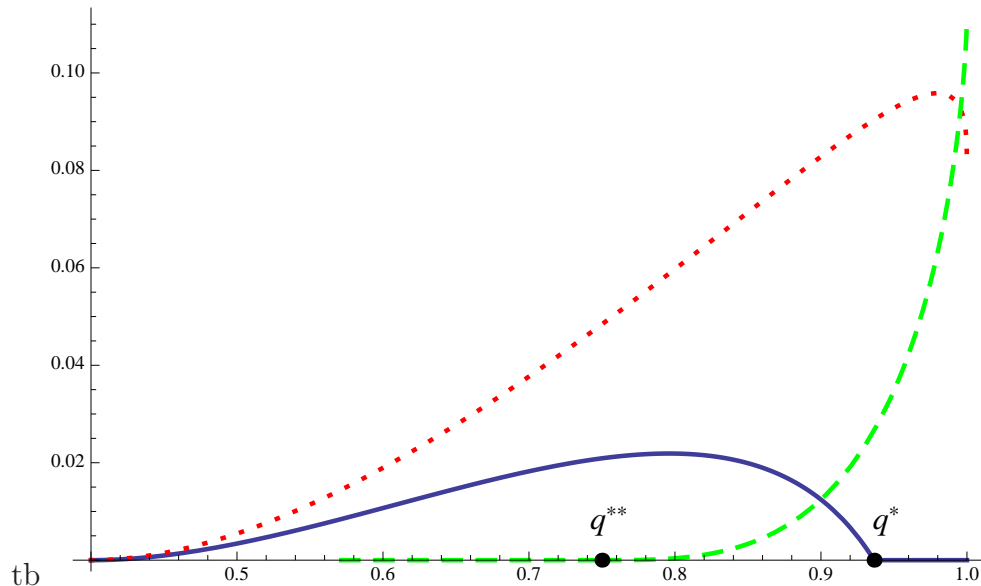


Figure 1: The principal's payoff (vertical axis) from the recursive Markov equilibrium, as a function of the probability q that the project is good (horizontal axis). The parameters are $c/r = 1$ for all curves. For the dotted curve, $(\psi, \sigma) = (3, 27/10)$, giving a high surplus, impatient project, with no delay and the principal's value positive throughout. For the dashed curve, $(\psi, \sigma) = (3/2, 5/4)$, giving a low surplus, impatient project, with delay and a zero principal value below the value $q^{**} = 0.75$. For the solid curve, $(\psi, \sigma) = (3, 4)$, giving a high surplus, patient project, with delay and a zero principal value for $q > q^* \approx .94$. We omit the case of a low surplus, patient project, where the principal's payoff is 0 for all q .

3.2 Non-Markov Equilibria

We now characterize the set of all equilibrium payoffs of the game. That is, we drop the restriction to recursive Markov equilibrium, though we maintain the assumption that equilibrium actions are pure on the equilibrium path.

This requires, as usual, to first understand how severely players might be credibly punished for a deviation, and thus, what each player's lowest equilibrium payoff is.

Here again, the limiting behavior as $dt \rightarrow 0$ admits an intuitive description, except for one boundary condition, which requires a fine analysis of the game with inertia (see Lemma 3).

3.2.1 Lowest Equilibrium Payoffs

Low-Surplus Projects ($\psi < 2$). We first discuss the relatively straightforward case of a low-surplus project. In the corresponding unique recursive Markov equilibrium, there is delay for all beliefs that are low enough (i.e., for all values of $q \in I$), where

$$I = \begin{cases} [\underline{q}, q^{**}] & \text{Impatient project,} \\ [\underline{q}, 1] & \text{Patient project.} \end{cases}$$

For these values of q , the principal's equilibrium payoff is zero. This implies that, for these beliefs, there exists a trivial non-Markov equilibrium in which the principal offers no funding on the equilibrium path, and so both players get a zero payoff; if the principal deviates and makes an offer, players revert to the strategies of the recursive Markov equilibrium, and so the principal has no incentive to deviate. Let us refer to this equilibrium as the “full-stop equilibrium.” This implies that, at least for q in I , there exists an equilibrium that drives down the agent's payoff to 0.

We claim that zero is also the agent's lowest equilibrium payoff for beliefs $\bar{q} > q^{**}$, in case $I = [\underline{q}, q^{**}]$. For each $q \in (\underline{q}, q^{**}]$, we construct a family of candidate non-Markov “no-delay” equilibria. Each member of this family corresponds to a different initial prior $\hat{q} > q$, and the candidate equilibrium path calls for no delay from the prior \hat{q} until the posterior falls to q (with the agent indifferent between working and shirking throughout), at which point there is a switch to the full-stop equilibrium. We thus have one such candidate equilibrium for each pair q, \hat{q} (with $\hat{q} > q \in (\underline{q}, q^{**}]$). We can construct an equilibrium with such an outcome if and only if the principal's payoff is nonnegative along the equilibrium path. For any $q \in (\underline{q}, q^{**}]$, the principal's payoffs will be positive throughout the path if \hat{q} is sufficiently close to q . However, if \hat{q} is too much larger than q , the principal's payoff will typically become negative. Let $\hat{q}(q)$ be the smallest $\hat{q} > q$ with the property that our the candidate equilibrium gives the principal a zero payoff at $\hat{q}(q)$ (if there exists such a $\hat{q} \leq 1$, and otherwise $\hat{q}(q) = 1$). Note that it must be that $\lim_{q \rightarrow \underline{q}} \hat{q}(q) = \underline{q}$, since otherwise the payoff function $v(q)$ defined by (9) would not be negative for values of q close enough to \underline{q} (which it is, because $\psi < 2$). Let $\hat{I} := I \cup \hat{q}(I)$ be the union of I and the image $\hat{q}(I)$ of I under the map \hat{q} . Then \hat{I} is an interval and it has length strictly greater than I .

It then suffices to argue that $1 \in \hat{I}$. It is in turn sufficient to show that $\hat{q}(q^{**}) = 1$. The fact that $q^{**} < 1$ indicates that the recursive Markov equilibrium featuring no delay until the posterior drops to q^{**} (with the agent's incentive constraint binding throughout), and then featuring delay for all subsequent posteriors, gives the principal a nonnegative payoff for all $q \in [q^{**}, 1]$. Our equilibrium differs from the recursive Markov equilibrium in terminating with a full stop at q^{**} instead of continuing (with delay) until the posterior hits \underline{q} . This structural difference potentially has two effects on the principal's payoff. First, the principal loses the continuation payoff that follows posterior q^{**} in the recursive

Markov regime. This payoff is zero, since delay follows q^{**} , and hence this has no effect on the principal's payoff. Second, the agent also loses the continuation payoff following q^{**} , which in the agent's case is positive. This reduces the agent's continuation payoff at every higher posterior. However, reducing the agent's continuation payoff reduces the payoff to shirking, and in a no-delay outcome in which the agent's incentive constraint binds, this reduction also reduces the share the principal has to offer the agent, increasing the principal's payoff. (See Appendix A.6.2 for details.) Hence, the principal's payoff in the candidate equilibrium is surely nonnegative, or $\hat{q}(q^{**}) = 1$.

We have established:¹³

Lemma 1 *Fix $\psi < 2$. For every $q > \underline{q}$, there exists an equilibrium for which the principal and the agent's payoffs converge to 0 as $dt \rightarrow 0$.*¹⁴

High-Surplus Projects ($\psi > 2$). This case is considerably more involved, as the unique recursive Markov equilibrium features no delay for initial beliefs that are close enough to \underline{q} , i.e. for all beliefs in J , where

$$J = \begin{cases} [\underline{q}, 1] & \text{Impatient project,} \\ [\underline{q}, q^*) & \text{Patient project.} \end{cases}$$

We must consider the principal and the agent in turn.

Because the principal's payoff in the recursive Markov equilibrium is not zero, we can no longer construct a full-stop equilibrium. Can we find some other non-Markov equilibrium in which the principal's payoff would be zero, so that we can replicate the arguments from the previous case? The answer is no: there is no equilibrium that gives the principal a payoff lower than the recursive Markov equilibrium, at least as $dt \rightarrow 0$. (For fixed $dt > 0$, her payoff can be driven slightly below the recursive Markov equilibrium, by a vanishing margin that nonetheless plays a key role in the analysis below.) Intuitively, by successively making the offers associated with the recursive Markov equilibrium, the principal can secure this payoff. The details behind this intuition are non-trivial, because the principal cannot commit to this sequence of offers, and the agent's behavior, given such an offer, depends on his beliefs regarding future offers. So we must show that there are no beliefs he could entertain about future offers that could deter the principal from making such an offer. Appendix A.7.2 establishes that the limit inferior (as $dt \rightarrow 0$) of

¹³For any q , the lowest principal and agent payoffs converge pointwise to zero as $dt \rightarrow 0$, but we can obtain zero payoffs for fixed $dt > 0$ only on some interval of the form $(q, 1)$ with $q > \underline{q}$. In particular, if $q_{t+dt} < \underline{q} < q$ (see (3) for q_{t+dt}), then the agent's payoff given q is at least cdt .

¹⁴More formally, for all $q > \underline{q}$, and all $\varepsilon > 0$, there exists $dt > 0$ such that both the principal's and the agent's lowest equilibrium payoff is below ε on $(q, 1)$ whenever the interval between consecutive offers is no larger than dt . For high-surplus, patient projects, for $q \geq q^*$, the same arguments as above yield that there is a worst equilibrium with a zero payoff for both players.

the principal's equilibrium payoff over all equilibria is the limit payoff from the recursive Markov equilibrium in that case:

Lemma 2 *Fix $\psi > 2$. For all $q \in I$, the principal's lowest equilibrium payoff converges to the recursive Markov (no-delay) equilibrium payoff, as $dt \rightarrow 0$.*

Having determined the principal's lowest equilibrium payoff, we now turn to the agent's lowest equilibrium payoff. In such an equilibrium, it must be the case that the principal is getting her lowest equilibrium payoff herself (otherwise, we could simply increase delay, and threaten the principal with reversion to the recursive Markov equilibrium in case she deviates; this would yield a new equilibrium, with a lower payoff to the agent). Also, in such an equilibrium, the agent must be indifferent between accepting or rejecting offers (otherwise, by lowering the offer, we could construct an equilibrium with a lower payoff to the agent).

Therefore, we must identify the smallest payoff that the agent can get, subject to the principal getting her lowest equilibrium payoff, and the agent being indifferent between accepting and rejecting offers. This problem looks complicated but turns out to be remarkably tractable, as explained below and summarized in Lemma 3. In short, there exists such a smallest payoff. It is strictly below the agent's payoff from the recursive Markov equilibrium, but it is positive.

To derive this result, let us denote here by v_M, w_M the payoff functions in the recursive Markov equilibrium, and by s_M the corresponding share (each as a function of q). Our purpose, then, is to identify all other solutions (v, w, s, λ) to the differential equations characterizing such an equilibrium, insisting on $v = v_M$ throughout, with a particular interest in the solution giving the lowest possible value of $w(\bar{q})$. Rewriting the differential equations (2) and (7) in terms of (v_M, w, s, λ) , and taking into account the delay λ (we drop the argument of λ in what follows), we get

$$0 = qps\pi - c - (r\lambda + pq)v_M(q) - pq(1 - q)v'_M(q), \quad (16)$$

and

$$0 = qp\pi(1 - s) - pq(1 - q)w'(q) - (r\lambda + qp)w(q) \quad (17)$$

$$= c - r\lambda w(q) - pq(1 - q)w'(q) + p(1 - q)w(q). \quad (18)$$

Since s_M solves (16) for $\lambda = 1$, any alternative solution (w, s, λ) with $\lambda > 1$ must satisfy (by subtracting (16) for $(s_M, 1)$ from (16) for (s, λ))

$$(\lambda - 1)v_M(q) = qp\pi(s - s_M).$$

Therefore, as is intuitive, $s > s_M$ if and only if $\lambda > 1$: delay enables the principal to increase her share. We can solve (17)–(18) for the agent's payoff given the share s and

then substitute into the (16) to get

$$r\lambda = \frac{qp\pi - 2c - pq(1-q)v'_M(q) - pw(q)}{v_M(q)} - pq.$$

Inserting into (18) and rearranging yields

$$q(1-q)v_M(q)w'(q) = w^2(q) + \left(v_M(q) + q(1-q)v'_M(q) + \frac{2 - (\psi + 2)q}{\sigma} \right) w(q) + \frac{v_M(q)}{\sigma}. \quad (19)$$

Because a particular solution to this Riccati equation (namely w_M) is known, the equation can be solved.¹⁵ Here, this gives that the general solution is

$$w(q) = w_M(q) + \frac{v_M(q)}{\xi(q)}, \quad (20)$$

where

$$-1 = q(1-q)\xi'(q) + \left(1 + \frac{\frac{2 - (\psi + 2)q}{\sigma} + 2w_M(q)}{v_M(q)} \right) \xi(q). \quad (21)$$

The factor $q(1-q)$ suggests working with the log-likelihood ratio $l = \ln \frac{q}{1-q}$. Considering (20), we might as well assume that $\xi(q) = 0$, for otherwise $w'(l) = w'_M(l)$, and we get back the known solution $w = w_M$. This gives the necessary boundary condition. We then obtain

$$w(q) = w_M(q) - v_M(q) \frac{\exp\left(\int_l^l h(v)dv\right)}{\int_l^l \exp\left(\int_l^v h(\chi)d\chi\right)dv}, \quad (22)$$

where l (and v) are log likelihood ratios, and where

$$h(l) := 1 + \frac{\left(2 - \frac{\psi + 2}{1 + e^{-l}}\right) \frac{1}{\sigma} + 2w_M(l)}{v_M(l)}.$$

To be clear, the Riccati equation admits two (and only two) solutions: w_M and w as given in (22). Let us study this latter function in more detail. We now use the expansions

$$v_M(q) = \frac{(\psi - 2)(\psi + 2)^3}{8\psi^2\sigma} (q - \underline{q})^2 \frac{c}{r} + O(q - \underline{q})^3, \quad w_M(q) = \frac{(\psi + 2)^2}{2\psi\sigma} (q - \underline{q}) \frac{c}{r} + O(q - \underline{q})^2.$$

¹⁵It is well-known that, given a Riccati equation

$$w'(q) = Q_0(q) + Q_1(q)w(q) + Q_2(q)w^2(q),$$

for which a particular solution w_M is known, the general solution is $w(q) = w_M(q) + g^{-1}(q)$, where g solves

$$g'(q) + (Q_1(q) + 2Q_2(q)w_M(q))g(q) = -Q_2(q).$$

We obtain (20)–(21) by applying this formula, and then changing variables to $\xi(q) := v_M(q)g(q)$.

This confirms an earlier observation: as the belief gets close to \underline{q} , payoffs go to zero, but they do so for two reasons for the principal: the maximum duration of future interactions vanishes (a fact that also applies to the agent), and the mark-up goes to zero. Hence, the principal's payoff is of the second order, while the agent's is only of the first. Expanding $h(l)$, we can then solve for the slope of w at \underline{q} , namely¹⁶

$$w'(\underline{q}) = \frac{\psi^2 - 4}{4\sigma\psi}.$$

Recall that $\psi > 2$, so that $0 < w'(\underline{q}) < w'_M(\underline{q})$. Again, the factor $(\psi - 2)$ should not come as a surprise: as $\psi \rightarrow 2$, the profit of the principal decreases, allowing her to credibly delay funding, and reduce the agent's payoff to compensate for this delay (so that her profit remains constant); when $\psi = 2$, her profit is literally zero, and she can drive the agent's payoff to zero as well.

This derivation is suggestive of what happens in the game with inertia, as the inertial friction vanishes, but it remains to prove that the candidate equilibrium $w < w_M$ is selected (rather than w_M).¹⁷ Appendix C.1 proves:

Lemma 3 *When $\psi > 2$ and $q \in I$, the infimum over the agent's equilibrium payoffs converges (pointwise) to \underline{w} , as given by (22), as $dt \rightarrow 0$.*

This solution satisfies $\underline{w}(q) < w_M(q)$ for all relevant $q < 1$: the agent can get a lower payoff than in the recursive Markov equilibrium. Note that, since the principal also obtains her lowest equilibrium payoff, it makes sense to refer to this payoff as the worst equilibrium payoff in this case as well.

Two features of this solution are noteworthy. First, it is straightforward to verify that for a high-surplus, patient project, i.e. when this derivation only holds for beliefs below q^* , the delay associated with this worst equilibrium grows without bound as $q \uparrow q^*$, and so the agent's lowest payoff tends to 0 (as does the principal's, by definition of q^*). This means that the worst equilibrium payoff is continuous at q^* , since we already know that it gives both players a payoff of 0 above q^* .

¹⁶A simple expansion reveals that $h(l) = \frac{8}{\psi-2} \frac{1}{(l-\underline{l})} + O(1)$, and defining $y(l) := \xi(q)$,

$$-1 = y'(\underline{l}) + \frac{8}{\psi-2} \frac{1}{(l-\underline{l})} y'(\underline{l}) (l-\underline{l}) + O(l-\underline{l})^2, \text{ or } y'(\underline{l}) = -\frac{\psi-2}{\psi+6},$$

which, together with $y(\underline{l}) = 0$, implies that $y(l) = -\frac{\psi-2}{\psi+6} (l-\underline{l}) + O(l-\underline{l})^2$. Using the expansion for v_M , it follows that $\frac{v_M(l)}{y(l)} = -\frac{\psi+6}{2(\psi+2)\sigma} (l-\underline{l}) + O(l-\underline{l})^2$, and the result follows from using $w'_M(\underline{l}) = \sigma^{-1}$.

¹⁷This relies critically on the multiplicity of recursive Markov equilibrium payoffs in the game with $dt > 0$, and in particular, the existence of equilibria with slightly lower payoffs. While this multiplicity disappears as $dt \rightarrow 0$, it is precisely what allows delay to build up as we consider higher beliefs, in a way to generate a non-Markov equilibrium whose payoff converges to this lower value w .




Figure 2: Functions w_M (agent’s recursive Markov equilibrium payoff, upper curve) and \underline{w} (agent’s lowest equilibrium payoff, lower curve), for $\psi = 3, \sigma = 2$ (high surplus, impatient project).

Second, for a high-surplus, impatient project, as $q \uparrow 1$, the solution $\underline{w}(q)$ given by Lemma 3 tends to one of two values, depending on parameters. These limiting values are exactly those obtained for the model in case information is complete: $\bar{q} = 1$. That is, the set of equilibrium payoffs for uncertain projects converges to the equilibrium payoff set for $\bar{q} = 1$, discussed in Section 4.2.

Figure 2 shows w_M and \underline{w} for the case $\psi = 3, \sigma = 2$.

3.2.2 The Entire Set of (Limit) Equilibrium Payoffs

The previous section determined the worst equilibrium payoff $(\underline{v}, \underline{w})$ for both the principal and the agent, given any prior q . As mentioned, this worst equilibrium payoff is achieved simultaneously for both parties. When this lowest payoff to the principal is positive, it is higher than her “minmax” payoff: if the agent never worked, the principal could secure no more than zero. Nevertheless, unlike in repeated games, this lower min-max payoff cannot be approached in an equilibrium: because of the sequential nature of the extensive-form, the principal can take advantage of the sequential rationality of the agent’s strategy to secure \underline{v} .

It remains to describe the entire set of equilibrium payoffs. This description relies on the following two observations.

First, for a given promised payoff w of the agent, the highest equilibrium payoff to

the principal that is consistent with the agent getting w , if any, is obtained by front-loading effort as much as possible. That is, equilibrium must involve no delay for some time, and then revert to as much delay as is consistent with equilibrium. Hence, play switches from no delay to the worst equilibrium. Depending on the worst equilibrium, this might mean full stop (for instance, if $\psi < 2$, but also if $\psi > 2$ and the belief at which the switch occurs, which is determined by w , exceeds q^*), or it might mean reverting to experimentation with delay (in the remaining cases). The initial phase in which there is no delay might be nonempty even if the recursive Markov equilibrium requires delay throughout. In fact, if reversion occurs sufficiently early (w is sufficiently close to \underline{w}), it is always possible to start with no delay, no matter the parameters. Appendix C.2 proves:

Proposition 2 *Fix $\bar{q} < 1$ and w . The highest equilibrium payoff available to the principal and consistent with the agent receiving payoff w , if any, is a non-Markov equilibrium involving no delay until making a switch to the worst equilibrium for some belief $q > \underline{q}$.*

As a function of the agent's payoff, the upper boundary of the payoff set, which gives the corresponding maximum payoff to the principal, is a single-peaked function of the agent's payoff. Given the prior belief, this boundary starts at the payoff $(\underline{v}, \underline{w})$, and initially slopes upward in w as we increase the duration of the initial no-delay phase. To identify the other extreme point, consider first the case in which $(\underline{v}, \underline{w}) = (0, 0)$. This is precisely the case in which, if there were no delay throughout (until the belief reaches \underline{q}), the principal's payoff would be negative. Hence, this no-delay phase must stop before the posterior belief reaches \underline{q} , and its duration is just long enough for the principal's (ex ante) payoff to be zero. The boundary thus comes down to a zero payoff to the principal, and her maximum payoff is achieved by some intermediate duration. Consider now the case in which $\underline{v} > 0$ (and so also $\underline{w} > 0$). This occurs precisely when no delay throughout (i.e., until the posterior reaches \underline{q}) is consistent with the principal getting a positive payoff; indeed, she then gets precisely \underline{v} , by Lemma 2. This means that, in this case as well, the boundary comes down eventually, with the other extreme point yielding the same minimum payoff to the principal, who achieves her maximum payoff for some intermediate duration in this case as well.

The second observation is that payoffs below this upper boundary, but consistent with the principal getting at least her lowest equilibrium payoff, can be achieved in a very simple manner. Because introducing delay at the beginning of the game is equivalent to averaging the payoff obtained after this initial delay with a zero payoff vector, varying the length of this initial delay, and hence the selected payoff vector on this upper boundary, achieves any desired payoff. This provides us with the structure of a class of equilibrium outcomes that is sufficient to span all equilibrium payoffs. The following summary is then immediate:

Proposition 3 *Any equilibrium payoff can be achieved with an equilibrium whose outcome features:*

1. *An initial phase, during which the principal makes no offer;*
2. *An intermediate phase, featuring no delay;*
3. *A final phase, in which play reverts to the worst equilibrium.*

Of course, any one or two of these phases may be empty in some equilibria. Observable deviations (i.e. deviations by the principal) trigger reversion to the worst equilibrium, while unobservable deviations (by the agent) are followed by optimal play given the principal's strategy.

In the process of our discussion, we have also argued that the favorite equilibrium of the principal involves an initial phase of no delay that does not extend to the end of the game. If $(\underline{v}, \underline{w}) = 0$, the switching belief can be solved in closed-form. Not surprisingly, it does not coincide with the switching belief for the only recursive Markov equilibrium in which we start with no delay, and then switch to some delay (indeed, in this recursive Markov equilibrium, we revert to some delay, while in the best equilibrium for the principal, reversion is to the full-stop equilibrium). In fact, it follows from this discussion that, unless the Markov equilibrium specifies no delay until the belief reaches \underline{q} , the recursive Markov equilibrium is Pareto-dominated by some non-Markov equilibrium.

3.2.3 Non-Markov Equilibria: Lessons

What have we learned from studying non-Markov equilibria? In answering this question we concentrate on efficient non-Markov equilibria.

First, these equilibria have a simple structure. The principal front-loads the agent's effort, calling for experimentation at the maximal rate until switching to the worst available continuation payoff. Second, this worst continuation involves terminating the project in some cases (low-surplus projects), in which the principal effectively has commitment power, but involves simply downsizing in other cases (high-surplus projects). Third, in either case, the switch to the undesirable continuation equilibrium occurs before reaching the second-best termination boundary \underline{q} of the recursive Markov equilibria. Finally, despite this inefficiently early termination or downsizing, the best non-Markov equilibria is always better for the principal (than the recursive Markov equilibrium) and can be better for both players.

The dynamic agency cost lurks behind all of these features. The premature switch to the worst continuation payoff squanders surplus, but in the process reduces the cost of providing incentives to the agent at higher posteriors. This in turn allows the principal to operate the experimentation without delay at these higher posteriors, giving rise to a surplus-enhancing effect sufficiently strong as to give both players a higher payoff than the recursive Markov equilibrium.

3.3 Summary

We can summarize the findings that have emerged from our examination of equilibrium. First, in addition to the standard agency cost, a dynamic agency cost arises out of the repeated relationship. The higher the agent's continuation payoff, the higher the principal's cost of inducing effort. This dynamic agency cost may make the agent so expensive that the principal can earn a nonnegative payoff only by slowing the pace of experimentation in order to reduce the future value of the relationship.

Recursive Markov equilibria without delay exist in some cases, but in others the dynamic agency cost can be effectively managed only by building in either delay for optimistic beliefs, or delay for pessimistic beliefs, or both. In contrast, non-Markov equilibria have a consistent and simple structure. The efficient non-Markov equilibria all front-load the agent's effort. Such an equilibrium features a initial period without delay, after which a switch is made to the worst equilibrium possible. In some cases this worst equilibrium halts experimentation altogether, and in other cases it features the maximal delay one can muster. The principal's preferences are clear when dealing with non-Markov equilibria—she always prefers a higher likelihood that the project is good, eliminating the non-monotonicities of the Markov case.

The principal always reaps a higher payoff from the best non-Markov equilibrium than from a recursive Markov equilibrium. Moreover, non-Markov equilibria may make both agents better off than the recursive Markov equilibrium. It is not too surprising that the principal can gain from a non-Markov equilibrium. Front-loading effort on the strength of an impending switch to the worst (possibly null) equilibrium reduces the agent's future payoffs, and hence reduces the agent's current incentive cost. But the eventual switch appears to waste surplus, and it is less clear how this can make both players better off. The cases in which both players benefit from such front-loading are those in which the recursive Markov equilibrium features some delay. Front-loading effectively pushes the agent's effort forward, coupling more intense initial effort with the eventual switch to the undesirable equilibrium. It can be surplus-enhancing to move effort forward, allowing both players to benefit.

We view the non-Markov equilibria as better candidates for studying venture capital contracts, where the literature emphasizes a sequence of full-funding stages followed by either downsizing or (more typically) termination.

4 Comparisons

Our analysis is based on three modeling conventions: there is moral hazard in the underlying principal-agent problem, there is learning over the course of repetitions of this problem, and the principal cannot commit to future contract terms. This section offers comparisons that shed light on each of these structural factors, and then turns to a

comparison with Bergemann and Hege [2].

4.1 Observable Effort

We first compare our findings to the case in which the agent's effort choice is observable by the principal. The most striking point of this comparison is that the ability to observe the agent's effort choice can make the principal worse off.

We assume that the agent's effort choice is unverifiable, so that it cannot be contracted upon, since contractibility obviates the agency problem entirely and returns us to the first-best world of Section 2.3. However, information remains symmetric, as the beliefs of the agent and of the principal coincide at all times.

4.1.1 Markov Equilibria

We start with Markov equilibria. The state variable is the common belief that the project is good. As before, let $v(q)$ be the value to the principal, as a function of the belief, and let $w(q)$ be the agent's value. If effort is exerted at time t , payoffs of the agent and of the principal must be given by, to the second order,

$$w(q_t) = pq_t\pi(1 - s(q_t))dt + (1 - r\lambda(q_t)dt)(1 - pq_tdt)w(q_{t+dt}) \quad (23)$$

$$\geq cdt + (1 - r\lambda(q_t)dt)w(q_t), \quad (24)$$

$$v(q_t) = (pq_t\pi s(q_t) - c)dt + (1 - r\lambda(q_t)dt)(1 - pq_tdt)v(q_{t+dt}).$$

Note the difference with the observable case, apparent from the comparison of the incentive constraint given in (24) with that of (7): if the agent deviates when effort is observable, his continuation payoff is still given by $w(q_t)$, as the public belief has not changed.

We focus on the frictionless limit (again, the frictionless game is not well-defined, and what follows is the description of the unique limiting equilibrium outcome as the frictions vanish). The incentive constraint for the agent must bind in a Markov equilibrium, since otherwise the principal could increase her share while still eliciting effort. Using this equality and taking the limit $dt \rightarrow 0$ in the preceding expressions, we obtain

$$0 = pq\pi(1 - s(q)) - (r\lambda(q) + pq)w(q) - pq(1 - q)w'(q) = c - r\lambda(q)w(q), \quad (25)$$

and

$$0 = pq\pi s(q) - c - (r\lambda(q) + pq)v(q) - pq(1 - q)v'(q). \quad (26)$$

As before, it must be that $q \geq \underline{q} = 2c/(p\pi)$, for otherwise it is not possible to give at least a flow payoff of c to the agent (to create incentives), while securing a return c to the principal (to cover the principal's cost). We assume throughout that $\underline{q} < 1$.

As in the case with unobservable effort, we have two types of behavior from which to construct Markov equilibria:

- The principal earns a positive payoff ($v(q) > 0$), in which case there must be no delay ($\lambda(q) = 1$).
- The principal delays funding ($\lambda(q) > 1$), in which case the principal's payoff must be zero ($v(q) = 0$).

Suppose first that there is no delay, so that $\lambda(q) = 1$ identically over some interval. Then we must have $w(q) = c/r$, since it is always a best response for the agent to shirk to collect a payoff of c , and the attendant absence of belief revision ensures that the agent can do so forever. We can then solve for $s(q)$ from (25), and plugging into (26) gives that

$$pq\pi - (2 + pq)c - (r + pq)v(q) - pq(1 - q)v'(q) = 0,$$

which gives as general solution

$$v(q) = \left[q - 2 + q \frac{\psi + 1}{1 + \sigma} - K(1 - q) \left(\frac{1 - q}{q} \right)^{\frac{1}{\sigma}} \right] \frac{c}{r}, \quad (27)$$

for some constant K . Considering any maximum interval over which there is no delay, the payoff v must be zero at its lower end (either experimentation stops altogether, or delay starts there), so that v must be increasing for low enough beliefs within this interval. Yet $v' > 0$ only if $K < 0$, in which case v is convex, and so it is increasing and strictly positive for all higher beliefs. Therefore, the interval over which $\lambda = 1$ and hence there is no delay is either empty or of the type $[q^*, 1]$. We can rule out the case $q^* = \underline{q}$, because solving for K from $v(\underline{q}) = 0$ gives $v'(\underline{q}) < 0$. Therefore, it must be the case that $v = 0$ and there is delay for some non-empty interval $[\underline{q}, q^*]$.

Alternatively, suppose that $v = 0$ identically over some interval. Then also $v' = 0$, and so, from (26), $s(q) = c/(pq\pi)$. From (25), $\lambda(q) = c/w(q)$, and so, also from (25),

$$pqw(q) = pq\pi - 2c - pq(1 - q)w'(q),$$

whose solution is, given that $w(\underline{q}) = 0$,

$$w(q) = \left[q(2 + \psi) - 2 - 2(1 - q) \left(\ln \frac{q}{1 - q} + \ln \frac{\psi}{2} \right) \right] \frac{c}{r}. \quad (28)$$

It remains to determine q^* . Using value matching for the principal's payoff at q^* to solve for K , value matching for w at q^* then gives that

$$q^* = 1 - \frac{1}{1 - 2 \frac{W_{-1} \left(-\frac{\psi - \sigma}{\psi} e^{-1 - \frac{\sigma}{2}} \right)}{\psi - \sigma}}, \quad (29)$$

where W_{-1} is the negative branch of the Lambert function.¹⁸ It is then immediate that $q^* < 1$ if and only if $\psi > \sigma$. The following proposition summarizes this discussion. Appendix D presents the foundations for this result, including an analysis for the case in which $dt > 0$ and a consideration of the limit $dt \rightarrow 0$.

Proposition 4 *The Markov equilibrium is unique, and is characterized by a value $q^* > \underline{q}$. When $q \in [\underline{q}, q^*]$, equilibrium behavior features delay, with the agent’s payoff given by (28) and zero payoff for the principal. When $q \in [q^*, 1]$, the equilibrium behavior features no delay, with the principal’s payoff given by (27) and positive for all $q > q^*$.*

[4.1] *If $\psi > \sigma$ (an impatient project), then q^* is given by (29).*

[4.2] *If $\psi < \sigma$ (a patient project), $q^* = 1$, and hence there is delay for all posterior beliefs.*

4.1.2 Non-Markov Equilibria

Because the principal’s payoff in the Markov equilibrium is 0 for any belief $q \in [\underline{q}, q^*]$, she is willing to terminate experimentation at such a belief, giving the agent a zero payoff as well. This is the analogue of the familiar “full-stop” equilibrium in the unobservable effort case. If both players expect that the project will be stopped at some belief q on the equilibrium path, there will be an equilibrium outcome featuring delay, and hence a zero principal’s payoff, for all beliefs in some interval above this threshold, which allows us to “roll back” and conclude that, for all beliefs, there exists a full-stop equilibrium in which the project is immediately stopped. Therefore, the lowest equilibrium payoffs are $(0, 0)$.

The principal’s best equilibrium payoff is now clear. Unlike in the unobservable case, there is no rent that the agent can secure by diverting the funds: any such deviation is immediately punished, as the project is stopped. Therefore, it is best for the principal that experimentation takes place for all beliefs above \underline{q} , without any delay, while keeping the agent at his lowest equilibrium payoff, i.e. 0. That is, $\lambda(q) = 1$ for all $q \geq \underline{q}$, and literally (in the limit) $s(q) = 1$ as well, with v solving, for $q \geq \underline{q}$,

$$(r + pq)v(q) = pq\pi - pq(1 - q)v'(q),$$

as well as $v(\underline{q}) = 0$. That is, for $q \geq \underline{q}$,

$$v(q) = \left[\frac{2 + \psi}{1 + \sigma} q \left(1 - \left(\frac{2(1 - q)/q}{\psi} \right)^{1 + \frac{1}{\sigma}} \right) \right] \frac{c}{r}. \quad (30)$$

We summarize this discussion in the following proposition. Appendix D again presents foundations.

¹⁸The positive branch only admits a solution to the equation that is below \underline{q} .

Proposition 5 *The lowest equilibrium payoff of the agent is zero for all beliefs. The best equilibrium for the principal involves experimentation without delay until $q = \underline{q}$, and the principal pays no more than the cost of experimentation. This maximum payoff is given by (30).*

The surplus from this efficient equilibrium can be divided arbitrarily. Because the agent's effort is observed and the principal's lowest equilibrium payoff is zero, we can ensure that the agent rejects any out-of-equilibrium offers by assuming that acceptance leads to termination. As a result, we can construct equilibria in which the agent receives an arbitrarily large share of the surplus, with any attempt by the principal to induce effort more cheaply leading to termination. In this way, it is possible to specify that the entire surplus goes to the agent in equilibrium. This gives the entire Pareto-frontier of equilibrium payoffs, and the convex hull of this frontier along with the zero payoff vector gives the entire equilibrium payoff set.

4.1.3 Comparison

One's natural inclination is to think that it can only help the principal to observe the agent's effort. Indeed, one typically thinks of principal-agent problems as trivial when effort can be observed. In this case, the principal may prefer to not observe effort. We see this immediately in the ability to construct non-Markov equilibria that divide the surplus arbitrarily. The ability to observe the agent's effort may then be coupled with an equilibrium in which the principal earns nothing, with the principal's payoff bounded away from zero when effort cannot be observed.

This comparison does not depend on constructing non-Markov equilibria. For patient projects, the principal's Markov equilibrium payoff under observable effort is zero, while there are Markov equilibria (for high surplus projects) under unobservable effort featuring a positive principal payoff. For impatient projects, the principal's Markov equilibrium payoff under observable effort is zero for pessimistic expectations, while it is positive for a high surplus project with unobservable effort. In each case, the observability is harmful for the principal.

4.2 Good Projects: $\bar{q} = 1$

Now suppose there is no learning. We continue to assume the interaction is repeated, since otherwise we have the standard principal-agent problem of Holmström [18], but we take $\bar{q} = 1$, so that the project is known to be good. Appendix E provides the detailed calculations behind the following discussion.

We first consider Markov equilibria. When $\bar{q} = 1$, a Markov equilibrium generates a stationary outcome. The same actions are repeated indefinitely, until the game is halted by a success. We find two types of Markov equilibria. For impatient projects, identified

by $\psi > \sigma$, the only Markov equilibrium features no delay until a success is obtained. For patient projects, the Markov equilibrium entails continual delay, with experimentation proceeding but at an attenuated pace, again until a success occurs. The principal earns a positive payoff in the former case, and a zero payoff in the latter.

When $\bar{q} = 1$, the project is inherently stationary—a failure leaves the players facing precisely the situation with which they started. One might then expect that the set of equilibrium payoffs is exhausted by considering equilibria with stationary outcomes, even if these outcomes are enforced by punishing with nonstationary continuation equilibria. To determine whether this is the case here, we must split the class of impatient projects into two categories. We say that a project is very impatient if

$$\psi < 2\sigma + \sigma^2.$$

In this case, the Markov equilibrium is the unique equilibrium, whether Markov or not. However, if we have a moderately impatient project, or

$$\sigma < \psi < 2\sigma + \sigma^2,$$

then there are non-Markov equilibria with stationary outcomes that give the principal a higher payoff than the Markov equilibrium, and there are non-Markov equilibria with nonstationary outcomes that give yet higher payoffs. For patient projects, or $\psi < \sigma$, there are equilibria with nonstationary outcomes that give *both* players a higher payoff than the Markov equilibrium.

As is the case when $\bar{q} < 1$, a simple class of equilibria spans the boundary of the set of all (weak perfect Bayesian) equilibrium payoffs. An equilibrium in this class features no delay for some initial segment of time, after which play switches to the worst equilibrium. This worst equilibrium is the full-stop equilibrium in the case of patient projects, and is an equilibrium featuring delay and relatively low payoffs in the case of a moderately impatient (but not very impatient) project. This is a stark illustration of the front-loading observed in the case of projects not known to be good. When $\bar{q} < 1$, the extremal equilibria feature no delay until the posterior probability q has deteriorated sufficiently, at which point a switch occurs to the worst equilibrium. When $\bar{q} = 1$, play occurs without delay and without belief revision, until switching to the worst equilibrium.

The benefits of the looming switch to the worst equilibrium are reaped up front by the principal in the form of lower incentive costs. It is then no surprise that there exist nonstationary equilibria of this type that provide higher payoffs to the principal than the Markov equilibrium. How can making the agent cheaper by reducing his continuation payoffs make the agent better off? The key here is that the Markov equilibrium of a patient project features perpetual delay. The nonstationarity front-loads the agent's effort, coupling a period without delay with an eventual termination of experimentation. This allows efficiency gains from which both players can benefit.

4.3 Commitment

Suppose the principal could commit to her future behavior, i.e., could announce at the beginning of the game a binding sequence of offers to be made to the agent.

Arguments analogous to those establishing Proposition 2 can be used to show (with the proof in Appendix F.1):

Proposition 6 *Fix $\bar{q} < 1$ and w . The highest equilibrium payoff (with commitment) available to the principal and consistent with the agent receiving payoff w , if any, involves no delay, until terminating experimentation at some belief $q > \underline{q}$.*

In the case of a low-surplus project, there exists an equilibrium without commitment providing a zero payoff to the principal, and hence there exists a full-stop equilibrium. This gives us the ability to construct equilibria consisting of no delay until the posterior belief reaches an arbitrarily chosen threshold, at which point experimentation ceases. But then there exist equilibria in which this threshold duplicates the posterior at which the commitment solution ceases experimentation. As a result, commitment does not increase the maximum value to the principal—anything the principal can do with commitment, the principal can also do without.

In the case of high-surplus projects, commitment is valuable. The principal’s maximal no-commitment payoff is achieved by eliciting no delay until a threshold value of q is reached at which point the players switch to the worst equilibrium. The latter still exhibits effort from the agent, albeit with delay, and the principal would fare better *ex ante* from the commitment to terminate experimentation altogether.

4.4 Powerless Principals

Finally, we compare our equilibria to the outcome described in Bergemann and Hege [2] (cf. Section 1.2). The principal has all of the bargaining power in our interaction. The primary modeling difference between our analysis and that of Bergemann and Hege [2] is that their agent has all of the bargaining power, making an offer to the principal in each period. How do the two outcomes compare? Perhaps the most striking result is the there are cases in which the agent may prefer the principal have the bargaining power.

We compare “recursive Markov” equilibria. Figures 3 and 4 illustrate the outcomes for the two models. The qualitative features of these diagrams are similar, though the details differ. The principal earns a zero payoff in any recursive Markov equilibrium in which the agent makes offers, and so the principal can only gain from having the bargaining power.

When σ is relatively small, there are values of $\psi > \sigma$ and $\psi > 2$ for which there is never delay (Case 1) when the principal makes offers but delay for low beliefs (Case 3) when the agent makes offers. For relatively low beliefs, the equilibrium when the principal makes the offers thus features a larger total surplus (because there is no delay), but the principal

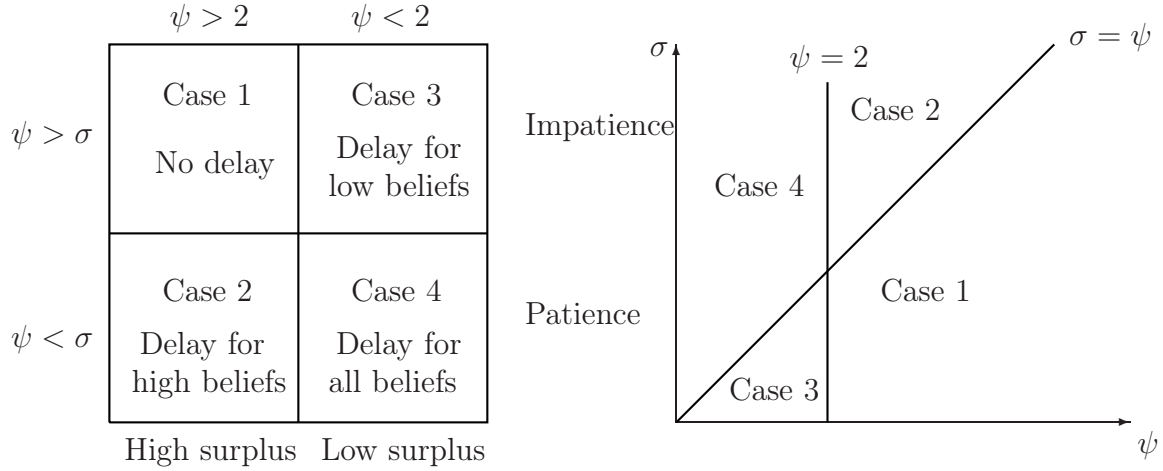


Figure 3: Illustration of Markov equilibria for the model considered in this paper, in which the principal makes offers.




Figure 4: Illustration of the outcomes described by Bergemann and Hege [2], in their model in which the agent makes offers. To represent the function $\sigma = 2 - 2p$, we fix $\pi/c =: \Pi$ and then use the definitions $\sigma = p/r$ and $\psi = p\Pi - 2$ to write $\sigma = 2 - 2p$ as $\sigma = 2 - 2(\psi + 2)/\Pi$. A straightforward calculation shows that the lines delineating the regions intersect at a value $\psi = \sigma > 2$.

captures some of this surplus. When the agent makes offers, delay reduces the surplus, but all of it goes to the agent. We have noted that when the principal makes offers, the proportion of the surplus going to the agent approaches one as q gets small. For relatively small values of q , the agent will then prefer that the principal has the bargaining power, allowing the agent to capture virtually all of a larger surplus.¹⁹

5 Summary

This paper makes three contributions: we develop techniques to work with repeated principal-agent problems in continuous time, we characterize the set of equilibria, and we explore the implications of institutional arrangements such as alternative allocations of the bargaining power. We concentrate here on the latter two contributions.

Our basic finding in Section 3 was that Markov and non-Markov equilibria differ significantly in both structure and payoffs. In terms of structure, the non-Markov equilibria that span the set of equilibrium payoffs share a simple common form, front-loading the agent's effort into a period of relentless experimentation followed by a switch to reduced or abandoned effort. This contrasts with Markov equilibria, which may call for either front-loaded or back-loaded effort.

Front-loading again plays a key role in the comparisons offered in Section 4. A principal endowed with commitment power would front-load effort, using her commitment in the case of a high-surplus project to increase her payoffs by accentuating this front-loading. When effort is observable, Markov equilibria either feature front-loading (impatient projects) or perpetual delay, while non-Markov equilibria eliminate delay entirely and achieve the first-best outcome. The principal may prefer (i.e., may earn a higher equilibrium payoff) being unable to observe the agent's effort, even if one restricts the ability to fortuitously select equilibria by concentrating on Markov equilibria. Similarly, the agent may be better off when the principal makes the offers (as here) than when the agent does.

References

- [1] Anat R. Admati and Paul Pfleiderer. Robust financial contracting and the role of venture capitalists. *Journal of Finance*, 49(2):371–402, 1994.
- [2] Dirk Bergemann and Ulrich Hege. The financing of innovation: Learning and stopping. *RAND Journal of Economics*, 36(4):719–752, 2005.

¹⁹Bergemann and Hege [2] show that in this case the scale of the project decreases as does q , bounding the ratio of the surplus under principal offers to the surplus under agent offers away from 1. This ensures that an arbitrarily large proportion of the former exceeds the latter.

- [3] Dirk Bergemann, Ulrich Hege, and Liang Peng. Venture capital and sequential investments. Cowles Foundation Discussion Paper 1682RR, Yale University, 2009.
- [4] James Bergin and W. Bentley MacLeod. Continuous time repeated games. *International Economic Review*, 34(1):21–37, 1993.
- [5] Donald A. Berry and Bert Fristedt. *Bandit Problems: Sequential Allocation of Experiments*. Springer, New York, 1985.
- [6] Bruno Biais, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet. Dynamic security design: Convergence to continuous time and asset pricing implications. *Review of Economic Studies*, 74(2):345–390, 2007.
- [7] Asher A. Blass and Oved Yosha. Financing R&D in mature companies: An empirical analysis. *Economics of Innovation and New Technology*, 12(5):425–448, 2003.
- [8] Patrick Bolton and Mathias Dewatripont. *Contract Theory*. MIT Press, Cambridge, Massachusetts, 2005.
- [9] Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.
- [10] Francesca Cornelli and Oved Yosha. Stage financing and the role of convertible securities. *Review of Economic Studies*, 70(1):1–32, 2003.
- [11] Diane K. Denis and Dilip K. Shome. An empirical investigation of corporate asset downsizing. *Journal of Corporate Finance*, 11(3):427–448, 2005.
- [12] Drew Fudenberg, David Levine, and Jean Tirole. Infinite-horizon models of bargaining with one-sided incomplete information. In Alvin E. Roth, editor, *Game-Theoretic Models of Bargaining*, pages 73–98. Cambridge University Press, Cambridge, 1983.
- [13] Dino Gerardi and Lucas Maestri. A principal-agent model of sequential testing. *Theoretical Economics*, 7(3):425–463, 2011.
- [14] Paul A. Gompers. Optimal investment, monitoring, and the staging of venture capital. *Journal of Finance*, 50(5):1461–1489, 1995.
- [15] Paul A. Gompers and Josh Lerner. The venture capital revolution. *Journal of Economic Perspectives*, 15(2):145–178, 2001.
- [16] Bronwyn H. Hall. The financing of innovation. In Scott Shane, editor, *Handbook of Technology and Innovation Management*, pages 409–430. Blackwell Publishers, Oxford, 2005.

- [17] Martin Hellwig. A note on the implementation of rational expectations equilibria. *Economics Letters*, 11:1–8, 1983.
- [18] Bengt Holmström. Moral hazard and observability. *Bell Journal of Economics*, 10(1):74–91, 1979.
- [19] Boyan Jovanovic and Balázs Szentes. On the return to venture capital. NBER working paper 12874, New York University and University of Chicago, 2007.
- [20] Steven N. Kaplan and Per Strömberg. Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. *Review of Economic Studies*, 70(2):281–315, 2003.
- [21] Godfrey Keller and Sven Rady. Strategic experimentation with Poisson bandits. *Theoretical Economics*, 5(2):275–311, 2010.
- [22] David Martimort and Jean-Jacques Laffont. *The Theory of Incentives: The Principal-Agent Model*. MIT Press, Cambridge, Massachusetts, 2002.
- [23] Robin Mason and Juuso Välimäki. Dynamic moral hazard and stopping. Department of Economics, University of Exeter and Helsinki School of Economics, 2011.
- [24] Carinne Peeters and Bruno van Pottelsberghe de la Potterie. Measuring innovation competencies and performances: A survey of large firms in Belgium. Institute of Innovation Research Working Paper 03-16, Hitotsubashi University, Japan, 2003.
- [25] Motty Perry and Philip J. Reny. A non-cooperative bargaining model with strategically timed offers. *Journal of Economic Theory*, 59:55–77, 1993.
- [26] Michael Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9:185–202, 1974.
- [27] William A. Sahlman. The structure and governance of venture-capital organizations. *Journal of Financial Economics*, 27(2):473–521, 1990.
- [28] Leo K. Simon and Maxwell B. Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(6):1171–1214, 1989.

A Online Appendix: Foundations

This section develops the foundations for results presented in Section 3. We assume there is a length of time $\Delta > 0$ such that if the principal has made an offer to the agent at time t , the principal cannot make another offer until time $t + \Delta$ (though the principal could wait longer to make the next offer). We characterize the set of equilibria for $\Delta > 0$ and then examine the limit as $\Delta \rightarrow 0$. The cost to be advanced to the agent is given by $c\Delta$ and the probability that a good project is a success is given by $p\Delta$. The discount factor is $\delta = e^{-r\Delta}$.

A.1 Outline

Section A.3 opens with a basic observation. Whether an agent facing an offer prefers to work or shirk depends not only on the usual suspects, such as the offer and the players' beliefs, but also on whether the agent is *expected* to work or shirk. We will find cases in which the agent will find it optimal to work if expected to do so, and to shirk if so expected. This will give rise to multiple equilibria. We will also find cases in which the agent will prefer to work if expected to shirk, and to shirk if expected to work. This will force us to work with mixed strategies (though we eventually show that these mixtures occur only off the equilibrium path), and precludes the existence of Markov equilibria. These observations color all of the subsequent analysis.

Section A.4 establishes some preliminary results that are important in formulating the problem in a manageable way, simplifying the types of beliefs we must consider and identifying the continuation payoffs for histories that will appear repeatedly in the analysis. Section A.5 introduces the solution concept of recursive Markov equilibrium, which we conventionally refer to simply as Markov equilibrium.

Section A.6 introduces an obvious candidate for a recursive Markov equilibrium, involving no delay, and establishes the conditions under which it is indeed an equilibrium. There are other Markov no-delay equilibria, and Section A.7 characterizes the set of such equilibria. Section A.8 characterizes the entire set of (recursive) Markov equilibria and establishes a limiting result. As $\Delta \rightarrow 0$, any sequence of Markov equilibria converges to the behavior considered in Section 3. It is this limiting uniqueness result that allows us to work in the convenient frictionless limit.

The extension of these results to non-Markov equilibria is for the most part straightforward, and is presented in Section 3.2, with details that are not immediately contained in the proofs in Sections C.1 and C.2.

A.2 Strategies

A history h_t at date t is a full description of the actions of the players from time zero up to, but not including, time t . It specifies for every $t' < t$ whether the principal made no offer (wait, w), or whether an offer $s \in [0, 1]$ was made, as well as the agent's action in that case (shirk, S , or work, W). Formally, a history at time t is a function $h_t : [0, t) \rightarrow \{[0, 1] \times \{S, W\}\} \cup \{w\}$ such that, for all $\tau < t$, if $h_t(\tau) \in [0, 1] \times \{S, W\}$, then $h_t(t') = w$ for all $t' \in (\tau, \tau + \Delta)$, $t' < t$. Implicit in this description is that all experiments were unsuccessful, as the game ends otherwise. A public history h_t^P is defined in the same way, but it only specifies the actions taken by the principal (i.e., it takes values in $[0, 1] \cup \{w\}$). The set of such functions are denoted H_t and H_t^P , and we write $h_{t|t'}^P$, $t' < t$, for the truncated history in $H_{t'}^P$ obtained from h_t^P (h_t^P is said to be a continuation of $h_{t|t'}^P$).

A (behavior) strategy for the principal is a collection $\sigma^P = (\sigma_t^P)_{t \in \mathbb{R}_+}$, where the σ_t^P are probability transitions from H_t^P into $[0, 1] \times \{w\}$, such that, for all t' and t with $t \in (t', t' + \Delta)$, we have $h_t^P(t') \in [0, 1]$ implies that $\sigma^P(h_t^P) = \{w\}$. The strategy of the agent σ^A is defined similarly (with H_t replacing H_t^P), given an outstanding offer s . To ensure that every pair σ of strategies uniquely determines a continuation-path, it is furthermore necessary to (innocuously) assume that, for all $t > t' \geq 0$, if $\sigma^P(h_{t|t'}^P) = w$, there exists $t'' > t'$ such that $\sigma^P(h_{t|t'}^P) = w$ for all $\tau \in (t', t'')$. See Perry and Reny [25] for details. As usual, we write $\sigma^P|_{h_t^P}$ (resp. $\sigma^A|_{h_t}$) for the continuation strategy induced by the given history.

A.3 Expectations and the Agent's Incentives

This section identifies the source of much of the technical difficulties, arising out of the interaction between expectations concerning the agent's actions and the agent's incentives to take those actions. Throughout, let

$$\varphi(q) := \frac{q(1-p)}{1-pq}. \quad (31)$$

Then $\varphi(q)$ is the posterior belief the project is good, given prior belief q and that the agent undertook an experiment that ended in a failure. The probability of a success is proportional to Δ , and so p should be replaced by $p\Delta$ in this expression, but we will suppress Δ whenever doing so causes no confusion.

In equilibrium, the principal and the agent share the same belief. Suppose we have reached a period in which the principal and the agent both attach probability q to the project being good. The principal offers share s . Will the agent work?

Let $W(\mathbf{1}_q, q)$ be the expected value to the agent of a continuation equilibrium in which the agent attaches probability q to the project being good, and the principal's

belief attaches a mass of probability one to belief q . If the equilibrium expectation is that the agent will work, then the agent will do so if

$$pq\pi(1-s) + (1-pq)\delta W(\mathbf{1}_{\varphi(q)}, \varphi(q)) \geq c + \delta W(\mathbf{1}_{\varphi(q)}, q).$$

The left side is the value of working, including the current expected payoff $pq\pi(1-s)$ from a success and the expected continuation payoff $(1-pq)\delta W(\mathbf{1}_{\varphi(q)}, \varphi(q))$ from a failure. The right side is the value of shirking, including the current payoff from expropriating c and the continuation payoff $\delta W(\mathbf{1}_{\varphi(q)}, q)$, with the agent now being more optimistic than the principal.

Let us suppose instead that the equilibrium expectation is that the agent will shirk. Then the incentive constraint is given by

$$pq\pi(1-s) + (1-pq)\delta W(\mathbf{1}_q, \varphi(q)) \leq c + \delta W(\mathbf{1}_q, q).$$

In each case, a larger agent share $1-s$ makes it more tempting for the agent to work. Hence, in the first case there is a cutoff s^W such that the agent will work for values $s < s^W$ and shirk for values $s > s^W$. In the second case, there is similarly a cutoff s^S .

Now let us consider three possibilities. First, it might be that for every history, $s^W = s^S$. In this case we can restrict attention to pure-strategy equilibria. The agent would work whenever s falls short of the value $s^W = s^S$ appropriate for the history in question, and shirk whenever s exceeds $s^W = s^S$. The principal's strategy would similarly be pure, solving an optimization problem subject to agent's incentive constraint. Finally, the Markov equilibrium would be unique in this case. The Markov assumption precludes constructing intertemporal incentives for the agent, and the principal would invariably extract as much surplus as possible from the agent, consistent with the agent still working, by setting $s = s^W = s^S$.

Appendix E, examining the case in which the project is known to be good ($\bar{q} = 1$) and so there is no learning, finds $s^W = s^S$. It is then no surprise that pure-strategy equilibria exist, and that Markov equilibria are unique. Similarly, we have $s^W = s^S$ in the case when actions are observable, examined in Section 4.1. The key in both cases is that the observed action (if any) and the outcome (failure) suffice to determine subsequent beliefs.

In contrast, when $\underline{q} < 1$ and actions are unobserved, subsequent beliefs depend on current expectations as to the agent's actions. Section A.6.4 shows that in this case, $s^W = s^S$ fails for almost all histories. Instead, a typical configuration is that $s^W > s^S$ for low values of q and $s^W < s^S$ for large values of q (though depending on parameters the latter interval may be empty).

Hence, and second, suppose we have a history at which $s^W > s^S$. For $s \in (s^S, s^W)$, the agent's optimal action depends on equilibrium expectations. The agent will prefer to work if expected to work, and prefer to shirk if expected to shirk. This allows us to construct multiple Markov equilibria, though the set of Markov equilibria converges to a unique limiting equilibrium as $\Delta \rightarrow 0$, described in Section 3.

Third, suppose we have a history at which $s^W < s^S$. For $s \in (s^W, s^S)$, the agent's optimal action again depends on equilibrium expectations. Here, however, the agent will prefer to shirk if expected to work, and will prefer to work if expected to shirk. This will preclude the construction of pure-strategy equilibria. Offers $s \in (s^S, s^W)$ will occur only off the equilibrium path, so that we can still restrict attention to equilibria featuring pure outcomes, but a complete specification of strategies (which is necessary to check that the principal finds such offers unprofitable, verifying that such offers are indeed off-path) will require mixing.

If the agent mixes, then not only will the principal and the agent subsequently have different beliefs (because only the agent knows the outcome of the mixture), but the principal's belief will attach positive probability to multiple agent beliefs. The support of the principal's belief q^P in any particular period will be a finite set, corresponding to the finitely many histories of actions the agent can have taken, but the maximum number of possible elements in this set grows with the passing of each period.

Appendix A.9 shows that cases in which $s^W < s^S$ arise, and cases in which $s^W > s^S$ can arise, when the agent makes the offers instead of the principal, as in Bergemann and Hege [2]. Hence, multiplicity and non-existence of Markov equilibria arise in that context as well.

A.4 Preliminaries

This section collects some results that simplify the principal's problem.

A.4.1 The Horizon is Effectively Finite

There is in principle no limit on the number of experiments the principal might induce the agent to conduct. However, there is an upper bound on the length of equilibrium experimentation. Appendix A.10 proves:

Lemma 4 *For any prior belief \bar{q} and waiting time Δ , there is a finite $T(\bar{q}, \Delta)$ such that there is no equilibrium attaching positive probability to an outcome in which the principal makes more than $T(\bar{q}, \Delta)$ offers to the agent.*

The intuition is straightforward. Every experiment pushes the posterior probability that the project is good downward, while costing at least c . There is then a limit on how many failures the principal will endure before becoming so pessimistic as to be unwilling to fund further experimentation.

The advantage of this result is that it makes available backward-induction arguments.

A.4.2 Two Beliefs are Enough

The current state of the project is described by a private belief q^A for the agent, and a public distribution q^P over beliefs for the principal. The public belief q^P potentially attaches positive probability to a finite number of posterior probabilities that the project is good, corresponding to the finite number of work/shirk combinations that the agent can have implemented in the preceding periods. However, we can restrict attention to rather simple instances of the public belief, attaching positive probability to at most two beliefs. Section A.11 proves:

Lemma 5 *Let $\{s_n\}_{n=1}^T$ be a sequence of offers, made by the principal at times $\{t_n\}_{n=1}^T$. Let the agent play a best response to this sequence of offers, and let q^P be the induced public belief. Then for sufficiently small Δ , after any initial subsequence of offers $\{s_n\}_{n=1}^t$ for $t \leq T$, the induced belief q^P attaches positive probability to at most two beliefs, given by q and $\varphi(q)$ for some q .*

Notice that we make no assumptions as to the nature of the offers $\{s_n\}_{n=1}^T$, and in particular do not require these to be part of an equilibrium outcome. We use here only the assumptions that the agent is playing a best response, and that the principal forms expectations correctly. We rely on Lemma 4 in restricting attention to a finite sequence of offers, which allows a backward-induction proof.

The key to proving Lemma 5 is to show that whenever an agent holding belief q^A is willing to work, any agent holding the more optimistic belief $\tilde{q}^A > q^A$ must strictly prefer to work. A more optimistic agent views a success as being more likely, and hence has more to gain from working, making it intuitive that an optimistic agent prefers to work whenever a pessimistic agent does. However, the result is not completely straightforward, since a more optimistic agent also faces a brighter future following a failure, enhancing the value of shirking. The proof verifies that the former effect is the more powerful.

A.4.3 The Value of an Optimistic Agent

Consider a special class of candidate equilibria, those in which the agent responds to every offer along the equilibrium path by working. In the context of such an equilibrium, the principal's belief after any history in which the principal has made no deviations will be of the form $\mathbf{1}_q$ for some q . The agent's belief will duplicate q if the agent has similarly made no deviations. If the agent instead shirks at least once, then the agent will hold a different belief, say \tilde{q} . Section A.12 proves:

Lemma 6 *In any equilibrium in which the agent is willing to work along the equilibrium path,*

$$\forall \tilde{q} \geq q : W(\mathbf{1}_q, \tilde{q}) = \frac{\tilde{q}}{q} W(\mathbf{1}_q, q). \quad (32)$$

The intuition for this result comes in two parts. First, the candidate equilibrium calls for the agent to always work. An agent who is more optimistic about success than the principal ($\tilde{q} > q$) will be all the more anxious to work, as we have seen in Lemma 5, and hence the agent’s out-of-equilibrium behavior duplicates his equilibrium behavior.

Second, the agent’s higher beliefs then simply scale up all the success probabilities involved in the agent’s expected payoff calculation, leading to the linear relationship given by (32).

A.4.4 The Value of a Pessimistic Agent

We might expect the principal to offer the agent the minimal amount required to induce the agent to work. Hence, a natural candidate for equilibrium behavior is that in which not only does the agent respond to every offer along the equilibrium path by working, but in each case is indifferent between working and shirking. In this context, we can identify the value of a pessimistic agent. Section A.13 proves:

Lemma 7 *In any equilibrium in which the agent is expected to work and is indifferent between working and shirking along the equilibrium path,*

$$W(\mathbf{1}_q, \varphi(q)) = c + \delta W(\mathbf{1}_{\varphi(q)}, \varphi(q)).$$

An agent characterized by $\varphi(q)$ is “one failure more pessimistic” than an agent or principal characterized by belief q or $\mathbf{1}_q$. The implication is that such an agent shirks at the next opportunity, at which point the agent and principal’s beliefs are aligned, giving continuation value $W(\mathbf{1}_{\varphi(q)}, \varphi(q))$.

A.5 The Equilibrium Concept: Recursive Markov Equilibria

We begin by defining an intuitive special case of the (weak perfect Bayesian satisfying the no-signaling-what-you-don’t-know restriction) equilibrium, *Markov equilibria*. For a fixed pair of strategies, we say that an offer is *serious* if it is prescribed by the principal’s equilibrium strategy, and induces the agent to work with positive probability (according to the agent’s equilibrium strategy). These are the offers that lead to a revision of the principal’s belief.

The prescribed actions in a Markov equilibrium depend only on the posterior beliefs of the agent and the principal, as well as the delay since the last serious offer. More precisely, the strategy of the principal depends on the public posterior belief only—the distribution of beliefs that she entertains about the agent’s private belief, derived via Bayes’ rule from the public history of offers and the equilibrium strategies—and on the delay since the last serious offer. The equilibria we consider are such that the principal makes another offer

(accepted if the agent follows his equilibrium strategy) if and only if this delay exceeds $\Lambda(q^P)\Delta$ for some $\Lambda(q^P) \geq 1$ (thus, Λ is part of the description of the strategy).

The agent’s strategy depends on this public belief, on his private belief that the project is good (derived from the public history of offers and the agent’s private history of effort choices), on the outstanding offer, and on the delay since the last serious offer. Public and private beliefs coincide along the equilibrium path, but not necessarily off-path. Both beliefs are relevant in determining the agent’s behavior and payoff—to identify the agent’s optimal action, we must determine his payoff from deviating, at which point the beliefs differ.

We call such equilibria *Markov equilibria*. One might think of restricting strategies in a Markov equilibrium still further, allowing the principal’s actions to depend only on the public belief q^P and the agent’s actions to depend only this public belief, on his private belief q^A , and on the outstanding offer. In contrast, we have added one element of non-stationarity—the dependence of the principal’s strategy on the delay since the last serious offer. This is necessary if we are to think of our inertialess continuous-time game as the limit of its counterparts with inertia. In particular, when $\Delta > 0$, the principal (when able to make an offer) could conduct a private randomization between making an offer and waiting Δ . This introduces an expected delay in the time until the principal makes her next offer, even while her strategy depends only on the public belief. The deterministic delay in our inertialess limit is the limiting counterpart of this expected delay. Our restriction to equilibria featuring pure strategies on the equilibrium path thus allows us to capture the limits of mixed equilibria from the corresponding games with inertia.

Formally, a Markov equilibrium is an equilibrium $\sigma = (\sigma^P, \sigma^A)$ in which, (i) for all h_t^P and $h_{t'}^P$ such that $t - \sup \{ \tau : h_t^P(\tau) \in [0, 1] \} = t' - \sup \{ \tau : h_{t'}^P(\tau) \in [0, 1] \}$ and $q^P(h_t^P) = q^P(h_{t'}^P)$, it holds that $\sigma^P(h_t^P) = \sigma^P(h_{t'}^P)$, and (ii) for all $h_t, h_{t'}$ such that $q^P(h_t^P) = q^P(h_{t'}^P)$ and $q^A(h_t^P) = q^A(h_{t'}^P)$, it holds that $\sigma^P(h_t, s) = \sigma^P(h_{t'}^P, s)$ for all outstanding offers $s \in [0, 1]$.

Unfortunately, Markov equilibria do not exist for all parameters, as the earlier discussion in Section A.3 foreshadowed. This is a common feature of extensive-form games of incomplete information (see for instance Fudenberg, Levine and Tirole [12] and Hellwig [17]). The problem is that, for some “knife-edge” beliefs, there exist multiple Markov equilibria. These beliefs, however, are endogenous, since they depend on earlier decisions by players, and in turn these decisions depend on the specific Markov equilibrium that is being selected at the later stage, so that the latter play must “remember” the earlier decisions to select the appropriate continuation equilibrium. In bargaining games, it suffices to include the last offer to recover existence (giving the so-called weak Markov equilibria). Here, this is not enough.

We are accordingly led to the following recursive definition: a *recursive Markov equilibrium* is a strategy profile σ such that:

- (i) If, given h_t , Markov equilibria exist, then $\sigma|_{h_t}$ is a Markov equilibrium.
- (ii) If, given h_t , there exists no Markov equilibrium, and $h_{t'}$ is a continuation of h_t , then
 - if $(q^P(h_t^P), q^A(h_t)) = (q^P(h_{t'}^P), q^A(h_{t'}))$ (and, for the principal’s strategy, the delay since the last serious offer is the same after both histories), then $\sigma|_{h_t} = \sigma|_{h_{t'}}$;
 - if instead $(q^P(h_t^P), q^A(h_t)) \neq (q^P(h_{t'}^P), q^A(h_{t'}))$, then $\sigma|_{h_{t'}}$ must be a recursive Markov equilibrium.

In words, the continuation equilibrium must be Markov whenever possible; if beliefs do not change, continuation strategies remain the same; and if beliefs do change, the continuation strategy must be a recursive Markov equilibrium.

Recursive Markov equilibria in which there is no randomization on the equilibrium path are well-defined because (as we show in Section A.6) Markov equilibria exist when the public belief is low enough, from which we can work backward to construct recursive Markov equilibria.²⁰ By definition, recursive Markov equilibria coincide with Markov equilibria whenever those exist, and it is not hard to see that our definition coincides with weak Markov equilibrium in games in which those exist. We hereafter typically refer to a recursive Markov equilibrium in which there is no randomization on the equilibrium path simply as a Markov equilibrium. This is the class of equilibria that we shall characterize, and whose outcomes converge to a unique limit as $\Delta \rightarrow 0$.

A.6 A Candidate Markov Equilibrium: No Delay Principal Optimum

We begin by considering a candidate Markov equilibrium. The principal makes an offer to the agent immediately upon the expiration of each waiting period Δ since the previous offer, until the posterior falls below a threshold (in the event of continued failure), after which no further experimentation occurs. The agent is indifferent between working and shirking in each period, and responds to each offer by working. We refer to this as a no-delay equilibrium, since there is no feasible way to make offers more rapidly. Technically, these strategies feature $\Lambda(q) = 1$ for all q . We will see in Section A.7 that there may be multiple no-delay Markov equilibria, but that the one introduced here maximizes the principal’s payoff over the set of such equilibria.

²⁰The restriction to strategies in which there is no randomization on the equilibrium path ensures that we can make the backward induction on degenerate public beliefs.

A.6.1 The Strategies

We let q_1 denote the final belief at which the principal makes a serious offer to the agent. We then number beliefs and offers backwards from 1. From Bayes' rule, we have $q_{\tau-1} = \varphi(q_\tau)$.

The principal's offer s_τ at time τ must suffice to induce effort on the part of the agent, and hence must satisfy

$$pq_\tau\pi(1 - s_\tau) + (1 - pq_\tau)\delta W(\mathbf{1}_{q_{\tau-1}}, q_\tau) \geq c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_\tau) \quad (33)$$

$$= c + \delta \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \quad (34)$$

We assume in this candidate equilibrium that the principal invariably offers a share s_τ causing the incentive constraint (33)–(34) to hold with equality (returning to this assumption in Section A.7). In the last period, facing a public and private belief concentrated on q_1 and share s_1 , the agent's incentive constraint is

$$pq_1\pi(1 - s_1) = c. \quad (35)$$

Using (35) and then working backward via the equality versions of (33)–(34), we have defined the on-path portion of our strategies for the candidate full-effort equilibrium.

A.6.2 The Costs of Agency

If our candidate strategies are to be an equilibrium, they must generate a nonnegative payoff for the principal. The principal's payoff in the final period (in which experimentation takes place) is $pq_1\pi s_1 - c$. Using the incentive constraint (35), this is nonnegative only if $pq_1\pi - 2c \geq 0$. We can thus identify the failure boundary \underline{q} , with the property that the principal makes serious offers to the agent if and only if

$$q \geq \underline{q} = \frac{2c}{p\pi}. \quad (36)$$

Combining (31) with (34), we can write the agent's incentive constraint as

$$pq_\tau\pi(1 - s_\tau) - c \geq p \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \quad (37)$$

The principal's share must at least cover the cost of her expenditure c , or $pq_\tau\pi s_\tau \geq c$. Combining with (37), our proposed strategies are an equilibrium if and only if

$$pq_\tau\pi - 2c \geq p \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \quad (38)$$

The key observation in (38) is that as the agent's continuation value becomes more lucrative, it becomes more expensive to provide incentives for the agent. Experimenting exposes the agent to the risk that the project may be a success now, eliminating future returns. Shirking now ensures an immediate payment of c (the diverted funds) plus the prospect of future experimentation. The more lucrative the future, the more tempting it is to shirk, and hence the more expensive it is to induce effort. We thus have a "dynamic agency cost." From (36), a principal contracting with a myopic agent (i.e., an agent who stubbornly persists in taking his future payoff to be zero) would induce full effort from the agent until hitting the failure boundary $\underline{q} = 2c/p\pi$. The agent's recognition of a valuable future increases the cost of such effort, potentially making it impossible to sustain.

A.6.3 Positive Principal Payoffs?

Condition (38) may fail for some values $q_\tau > \underline{q}$, in which case there is no way to satisfy the incentive constraint given by (33) and still cover the principal's experimentation cost. Under these circumstances, our candidate strategies do not describe equilibrium behavior.

To identify conditions under which the principal's payoff is positive, we can rearrange the Bayesian updating expression given by (31) to obtain

$$\frac{1}{q_\tau} = 1 + \frac{1 - q_1}{q_1}(1 - p)^{\tau-1}. \quad (39)$$

To conserve on notation, let w_τ be the agent's payoff when facing offer τ of the candidate equilibrium. Then let us introduce the variable $\omega_\tau := w_\tau/(q_\tau c)$. Using the incentive constraint (34), the agent's payoff in the candidate Markov equilibrium solves

$$\omega_{\tau+1} = 1 + Q_1\beta^\tau + \delta\omega_\tau,$$

where $\beta = 1 - p$ and $Q_1 = (1 - q_1)/q_1$. This elementary difference equation has as solution

$$\omega_\tau = \frac{1 - (\omega_1 + 1)\delta^\tau}{1 - \delta} + \beta Q_1 \frac{\delta^\tau - \beta^\tau}{\delta - \beta} - \frac{\delta^{\tau-1}((1 - \delta)(Q_1\beta + 1) - \omega_1)}{1 - \delta}.$$

Similarly, we can let v_τ be the principal's payoff when making offer τ of the candidate equilibrium, define $\nu_\tau := v_\tau/(q_\tau c)$, and obtain

$$\nu_{\tau+1} = p\pi/c - 2(1 + Q_1\beta^{\tau+1}) + \delta[(1 - p)\nu_\tau - p\omega_\tau],$$

and so, after some algebra,

$$\nu_\tau = \frac{1 - \beta^{\tau-1}\delta^{\tau-1}}{1 - \beta\delta}(\psi+1) - \beta^2 Q_1 \frac{\delta^{\tau-1} - \beta^{\tau-1}}{\delta - \beta} - \frac{1 - \delta^{\tau-1}}{1 - \delta}(1 + \beta^\tau Q_1) - \delta^{\tau-1}\omega_1 + \beta^{\tau-1}\delta^{\tau-1}(\nu_1 + \omega_1), \quad (40)$$

(recalling that $\psi := p\pi/c - 2$).

The term ν_τ can be written as a weighted sum of $\delta^{\tau-1}$, $\beta^{\tau-1}$ and $(\delta\beta)^{\tau-1}$. It follows that the second difference $\nu_{\tau+2} + \nu_\tau - 2\nu_{\tau+1}$ can be written as a weighted sum of $\delta^{\tau-3}$, $\beta^{\tau-3}$ and $(\delta\beta)^{\tau-3}$, and hence these second differences can change sign (as τ varies) at most twice. Equivalently, in the limit as $\Delta \rightarrow 0$ the second derivative of the function ν_τ can change sign at most twice.²¹ It is also immediate that $q_1 p \pi - 2c = \nu_1 > \nu_0 = 0$, so if we let (noting that the set on the right may be empty, in which case we take its infimum to be ∞)

$$\tilde{\tau} = \inf \{ \tau > 1 : \nu_\tau < 0 \},$$

then it follows that ν must be concave for some $\tau \in \{0, \dots, \tilde{\tau}\}$, unless $\tilde{\tau} = \infty$. It follows that, considering increasing values of τ above $\tilde{\tau}$, either ν is concave throughout, or convex and then concave, or concave, convex and then concave. Let $\hat{\tau}$ denote the lowest $\tau > \tilde{\tau}$ for which $\nu_\tau > 0$,

$$\hat{\tau} = \min \{ \tau > \tilde{\tau} : \nu_\tau > 0 \},$$

if there exists such a τ , and otherwise take $\hat{\tau}$ to be ∞ , i.e., Then ν must have been convex for some value of $k \in \{\tilde{\tau}, \dots, \hat{\tau}\}$, and so above $\hat{\tau}$, the sequence ν is at most first positive, then negative. This argument establishes:

Lemma 8 *The value ν_τ solving (40) is positive for low values of τ , then (possibly) negative, then (possibly) positive, then (possibly) negative.*

Lemma 8 identifies some possibilities for ν_τ . We can specify more precisely the circumstances in which these various possibilities obtain by examining the limiting case of small Δ . Section A.14 proves:

Lemma 9

(9.1) *The (pointwise) limit of ν (as $\Delta \rightarrow 0$) has at most one inflection point, and the limit is positive for q close to, but above \underline{q} if $\psi > 2$ but not if $\psi < 2$. Hence, if $\psi < 2$, $q_{\tilde{\tau}} \rightarrow \underline{q}$.*

(9.2) *The (pointwise) limit of ν (as $\Delta \rightarrow 0$) admits a root $q^* \in (\underline{q}, 1)$ when $\psi > 2$ and $\sigma > \psi$.*

(9.3) *Because the second derivative of the limit of ν_τ at $q = \underline{q}$ is not zero, it is possible that $q_{\tilde{\tau}} \rightarrow \underline{q}$ as $\Delta \rightarrow 0$ (indeed, this does occur if $\psi < 2$), but then $q_{\hat{\tau}} \rightarrow \underline{q}$: if they exist, the first two intervals (i.e., $(\underline{q}, \tilde{q})$ and (\tilde{q}, \hat{q})) cannot both “vanish” in the limit.*

Combining Lemma (9.1) and Lemma (9.2) (and recalling that $\sigma := p/r$):

²¹These statements follow from the observation that the function $f(n) = ax^n + by^n + cz^n$, with parameters a, b, c, x, y, z , can have at most two zeros, which is a straightforward calculation.

Corollary 1 *For sufficiently small Δ , the no-delay principal-optimum strategies yield positive principal payoffs, and hence potentially generate an equilibrium outcome, if $\psi > 2$ and $\psi > \sigma$, but do not constitute an equilibrium if either $\psi < 2$ or $\psi < \sigma$.*

After acquiring a preliminary result in the next section, Section A.6.5 completes the specification of the strategies, showing in the process that we indeed have an equilibrium when ψ exceeds both 2 and σ , while Section A.8 characterizes Markov equilibria for other parameter values.

A.6.4 The Agent's Incentives

Let us assume that $\psi > 2$ and $\psi > \sigma$, so that our no-delay principal-optimum strategies are a candidate for equilibrium. Define $s^W(q)$ to be the value of s that solves the incentive constraint (34) with equality. Our candidate Markov equilibrium then calls for the principal to offer share $s^W(q)$ in every period, with the agent working in response to smaller values of s and shirking in response to larger values of s .

How does the agent respond to other values of s ? This depends on the relative magnitudes of s^W and s^S . Appendix A.15 proves the following:

Lemma 10 *There exists a value $\tilde{q}(\Delta) \geq \underline{q}$ such that*

$$\begin{aligned} s^S(q) &\leq s^W(q) & \text{if } q < \tilde{q}(\Delta), \\ s^S(q) &> s^W(q) & \text{if } q > \tilde{q}(\Delta). \end{aligned}$$

The value of $\tilde{q}(\Delta)$ remains bounded away from \underline{q} as $\Delta \rightarrow 0$. There exist parameter values for which $\tilde{q}(\Delta) < 1$, and remains bounded below 1 as $\Delta \rightarrow 0$.

A.6.5 Completing the Strategies

We now specify the strategies in our no-delay principal-optimum Markov equilibrium.

Let us start with the slightly simpler case in which, given the public history, the principal's belief is degenerate. Suppose the players face posterior belief q_τ , where τ additional failures would give a posterior exceeding \underline{q} , while $\tau + 1$ additional failures would give a posterior falling short of \underline{q} . The principal's strategy is straightforward. Facing posterior q_τ , the principal makes offer $s^W(q_\tau)$. If we have $s^S(q_\tau) < s^W(q_\tau)$, then the agent's strategy is similarly straightforward: in each period, along the equilibrium path, the agent works if and only if $s \leq s^W(q_\tau)$.

Suppose $s^S(q_\tau) > s^W(q_\tau)$. Then we specify strategies as:

- The agent works if $s \leq s^W(q_\tau)$. Play continues with the principal offering $s^W(\varphi(q_\tau))$ next period.

- The agent shirks if $s \geq s^S(q_\tau)$. Play continues with the principal offering $s^W(q_\tau)$ next period.
- If $s \in (s^W(q_\tau), s^S(q_\tau))$, the agent mixes, with probability $\rho(s, q_\tau)$ of working. The principal then enters the next period with mass on two possible agent types. The principal induces both types to work with each of the next $z(s, q_\tau)$ offers, for some $z(s, q_\tau) \in \{0, \dots, \tau - 1\}$, in each case causing the subsequent period to be reached with q^P attaching positive probability to two agent beliefs. In the $z(s, q_\tau) + 1$ st period, the principal mixes between causing only the more optimistic agent to work and causing both to work (attaching nonzero but possibly unitary probability to the former). If the latter is the case, only the more optimistic agent is induced to work in period $z(s, q_\tau) + 2$. Thereafter the principal's belief attaches positive probability to only a single agent belief.

The first step in showing that this is an equilibrium is to characterize the mixture $\rho(s, q_\tau)$, the period $z(s, q_\tau)$, and the principal's mixture in that period. Sections A.16–A.17 prove:

Lemma 11 *There exist an agent mixture $\rho(s, q_\tau)$, a period $z(s, q_\tau)$, and a nonzero mixture with which only the optimistic agent is induced to work in period $z(s, q_\tau) + 1$, such that (i) the agent is indifferent between working and shirking in response to offer s , making the mixture $\rho(s, q_\tau)$ a best response for the agent, (ii) the principal prefers inducing only the optimistic agent to work in period $z(s) + 1$ or $z(s) + 2$ to doing so in any other period, and (iii) the principal either prefers to induce this outcome in period $z(s, q_\tau) + 1$ if the mixture in that period is unitary, or otherwise is indifferent between doing so in period $z(s, q_\tau) + 1$ and $z(s, q_\tau) + 2$.*

Next, we show that the result is a Markov equilibrium. Section A.18 proves:

Lemma 12 *Let the principal's belief be given by $\mathbf{1}_q$ for some q . For sufficiently small Δ , any offer $s \in (s^W(q), s^S(q))$ gives the principal a lower payoff than does $s^W(q)$.*

Let us now describe strategies *off* the equilibrium path. Given Lemma 5, we consider a public history h_t^P that gives rise to a pair of beliefs q and $\tilde{q} = \varphi(q)$, along with a probability μ attached to q . That is, the principal attaches probability μ to the agent having private belief q , and $1 - \mu$ to the (slightly more pessimistic) belief \tilde{q} . As in section A.3, we can associate to the beliefs q, \tilde{q} two thresholds s^W, \tilde{s}^W , and s^S, \tilde{s}^S . By Lemma 5, $\max\{\tilde{s}^W, \tilde{s}^S\} < \min\{s^W, s^S\}$. The principal offers either s^W or \tilde{s}^W , according to which is more profitable; if they are equally profitable, she randomizes between those two so as to vindicate the indifference between accepting and rejecting of the agent's type who randomized last along the history h_t^P .

Given an outstanding offer, there are four possibilities, depending on $\tilde{s}^W \gtrless \tilde{s}^S$, and $s^W \gtrless s^S$.

- if $\tilde{s}^W < \tilde{s}^S$, and $s^W < s^S$:
 1. if $s < \tilde{s}^W$, then both types work;
 2. if $s \in (\tilde{s}^W, \tilde{s}^S)$, then type \tilde{q} randomizes while type q works. The randomization is such that the principal is indifferent between having both types and only the optimistic type work in a later period in a way that allows her to randomize between the two so as to make type \tilde{q} indeed indifferent.
 3. if $s \in [\tilde{s}^S, s^W]$, type \tilde{q} shirks while type q works;
 4. if $s \in (s^W, s^S)$, then type q randomizes while type \tilde{q} shirks. The randomization is such that the principal is indifferent between having both types and only the optimistic type work in a later period in a way that allows her to randomize between the two so as to make type q indeed indifferent.
 5. if $s \geq s^S$, both types shirk.
- if $\tilde{s}^W < \tilde{s}^S$, yet $s^W \geq s^S$:
 1. if $s < \tilde{s}^W$, then both types work;
 2. if $s \in (\tilde{s}^W, \tilde{s}^S)$, then type \tilde{q} randomizes while type q works. The randomization is such that the principal is indifferent between having both types and only the optimistic type work in a later period in a way that allows her to randomize between the two so as to make type \tilde{q} indeed indifferent.
 3. if $s \in [\tilde{s}^S, s^W)$, type \tilde{q} shirks while type q works;
 4. if $s \geq s^W$, both types shirk.
- if $\tilde{s}^W \geq \tilde{s}^S$, yet $s^W < s^S$:
 1. if $s < \tilde{s}^W$, then both types work;
 2. if $s \in [\tilde{s}^W, s^W]$, then type \tilde{q} shirks while type q works;
 3. if $s \in (s^W, s^S)$, then type q randomizes while type \tilde{q} shirks. The randomization is such that the principal is indifferent between having both types and only the optimistic type work in a later period in a way that allows her to randomize between the two so as to make type q indeed indifferent;
 4. if $s \geq s^W$, then both types shirk.
- if $\tilde{s}^W \geq \tilde{s}^S$, and $s^W \geq s^S$:
 1. if $s < \tilde{s}^W$, then both types work;
 2. if $s \in [\tilde{s}^W, s^W)$, then type \tilde{q} shirks while type q works;
 3. if $s \geq s^W$, then both types shirk.

This completes the description of equilibrium strategies. We then have to check whether sequential rationality is satisfied off the equilibrium path. The key question here is whether the principal would find one of \tilde{s}^W or s^W optimal. Section A.19 proves:

Lemma 13 *Suppose q^P attaches probability to two beliefs, q and $\tilde{q} = \varphi(q)$. Then the optimal offer for the principal is one of \tilde{s}^W or s^W .*

A.6.6 Summary: No-Delay Principal-Optimum Markov Equilibrium

We have established:

Proposition 7 *Let $\psi > 2$ and $\psi > \sigma$. Then for sufficiently small Δ , there exists a Markov equilibrium in which, whenever $q > 2c/p\pi$, the principal makes an offer at every opportunity, each such offer makes the agent indifferent between working and shirking, and the agent works.*

A.7 Other No-Delay Markov Equilibria

We can identify a collection of additional no-delay Markov equilibria.

A.7.1 The Final Period

We begin by examining the final period, beginning with a posterior q_1 featuring

$$\varphi(q_1) < \underline{q} = \frac{2c}{p\pi} < q_1.$$

Hence, one more failed experiment will make the principal too pessimistic to continue. The payoffs in the no-delay principal-optimum Markov equilibrium are then

$$\begin{aligned} V(\mathbf{1}_{q_1}, q_1) &= pq_1\pi - 2c, \\ W(\mathbf{1}_{q_1}, q_1) &= c. \end{aligned}$$

These payoffs place an upper bound on the principal's payoff in a no-delay equilibrium, and a lower bound on the agent's payoff in a no-delay equilibrium. Let \tilde{q} be such that $\underline{q} = \varphi(\tilde{q})$. Section A.20 proves:

Lemma 14 *There exists $\hat{q} \in (\underline{q}, \tilde{q})$ such that for $q_1 \in [0, \hat{q}]$, the range of principal payoffs achievable in a no-delay Markov equilibrium (and in any no-delay equilibrium) is $[0, pq_1\pi - 2c]$. For $q_1 \in [\hat{q}, \tilde{q})$, the range is $[pq_1\pi - \frac{2-\delta p}{1-\delta p}c, pq_1\pi - 2c]$.*

For $q \in (\hat{q}, \tilde{q})$, we have $0 < pq_1\pi - \frac{2-\delta p}{1-\delta p}c < pq_1\pi - 2c$. Hence, we see that in the final period, there is a range of equilibrium payoffs for the principal. In addition, the principal's minimum equilibrium payoff is zero for some posteriors, but for some posteriors, the principal's payoff is strictly positive in the final period.

A.7.2 Constructing the Set of No-Delay Equilibria

We can work backwards from the final period to construct the set of no-delay Markov equilibria. In the course of doing so, beliefs will run through a set of posteriors $\{q_\tau\}_{\tau=1}^\infty$, which we can take as fixed throughout. We will generate a range of equilibrium payoffs in each period. There are potentially two degrees of freedom in constructing these equilibria that fix the upper and lower bounds of the range—the choice of continuation payoffs and the choice of current shares. To maximize the principal’s payoff, we choose the lowest equilibrium continuation payoff for the agent and choose share s^W . To minimize the principal’s payoff, we choose the largest continuation payoff for the agent and share s^S if $s^S < s^W$. The latter choice will be available for small values of q , but may not be available for large values of q . This procedure generates the entire set of no-delay Markov equilibria, as long as the principal’s payoff remains positive.

If the principal’s payoff is positive, then the multiplicity of the principal’s payoff disappears in the limit as $\Delta \rightarrow 0$. Section A.21 proves:

Lemma 15 *Let the no-delay principal-optimum strategies give the principal a positive payoff for posteriors in some interval $[\underline{q}, \tilde{q}]$ (and hence constitute an equilibrium for any $\bar{q} \in [\underline{q}, \tilde{q}]$). Then if $\bar{q} \in [\underline{q}, \tilde{q}]$, as $\Delta \rightarrow 0$, the lowest equilibrium payoff for the principal, over all equilibria, is positive and converges to the principal’s payoff from the no-delay principal-optimum Markov equilibrium.*

A.8 The Set of Markov Equilibria

Now we characterize the full set of Markov equilibria. The cases yet to be addressed are those in which either $\psi < 2$ or $\psi < \sigma$ holds.

A.8.1 A Canonical Equilibrium

We construct a canonical Markov equilibrium. We refer to an event in which the principal makes an offer as a period. The agent works in response to every offer, so for each posterior, the number of periods before the posterior crosses the termination threshold \underline{q} is fixed. The length of time between periods is at least Δ , but will be larger if there is delay. We work backward from period 1, the final period, as follows:

- (1) Let \bar{V}_τ be the largest principal payoff generated in period τ under a no-delay equilibrium. Let τ' be the first period, if any, in which $\bar{V}_{\tau'} < 0$. Then for $q \in [\underline{q}, q_{\tau'-1}]$, there exists a no-delay equilibrium with $V_{\tau'-1} = 0$ (see Lemma 16 below), and we choose this as our canonical equilibrium.
- (2) Working backwards from τ' , we insert just enough delay at each period τ to ensure that \bar{V}_τ , the largest payoff available to the principal at period τ , equals zero. We

then set s_τ to ensure $V_\tau = 0$. We continue to do this until (possibly) reaching a period τ'' at which, given $V_{\tau''-1} = 0$, a strictly positive principal payoff is available at period τ'' without delay.

- (3) Upon reaching such a period τ'' , we set $V_{\tau''-1} = 0$, and then work backwards constructing no-delay strategies.
- (4) This may continue until reaching a period τ''' in which no delay ensures $V_{\tau'''} < 0$. Then we choose the equilibria in periods $\{\tau'', \dots, \tau''' - 1\}$ so that $V_{\tau'''-1} = 0$, and once again work backward with delay to set V_τ thereafter equal to zero.

We are thus alternating between periods of no delay and positive principal payoffs and periods of delay and zero principal payoffs. Lemma 8 ensures that the regimes must come in the sequence described in our construction, and that we have identified the complete range of possibilities for such sequences.

If this procedure is to be well defined, we must show that whenever our no-delay principal-optimum construction reaches its first period τ with $V_\tau < 0$, there is an equilibrium with $V_{\tau-1} = 0$. Let \underline{V} denote the smallest no-delay Markov equilibrium payoff to the principal, and \overline{V} the largest such payoff. Section A.22 proves:

Lemma 16 *Fix a posterior q . Then, for sufficiently small Δ ,*

$$\underline{V}(\mathbf{1}_{\varphi(q)}, \varphi(q)) \leq \overline{V}(\mathbf{1}_q, q).$$

The implication of this is that the set of principal's payoffs for a given belief can never jump across zero (as we vary beliefs). If the smallest payoff for the principal at $\varphi(q)$ is positive, it cannot be that the largest payoff at q is negative.

A.8.2 Characterizing the Canonical Equilibrium

The canonical construction gives us periods of no delay and positive payoffs for the principal interspersed with periods of delay and zero payoffs for the principal. Section A.6.3 has characterized cases involving no delay and a positive payoff for the principal. Here, we examine delay.

The principal's payoff must be zero if there is to be delay, and hence

$$pq_{\tau+1}\pi s_{\tau+1} = c.$$

We then have

$$\begin{aligned} w_{\tau+1} &= pq_{\tau+1}(1 - s_{\tau+1})\pi - c + \delta(1 - p)\frac{q_{\tau+1}}{q_\tau}\Lambda_{\tau+1}w_\tau \\ &= c + \delta\Lambda_{\tau+1}\frac{q_{\tau+1}}{q_\tau}w_\tau. \end{aligned}$$

As before, let $\omega_\tau := w_\tau / (cq_\tau)$, giving

$$\omega_{\tau+1} = \frac{p\pi}{c} - \frac{1}{q_{\tau+1}} + \delta(1-p)\Lambda_{\tau+1}\omega_\tau = \frac{1}{q_{\tau+1}} + \delta\Lambda_{\tau+1}\omega_\tau,$$

and so

$$\delta\Lambda_{\tau+1}\omega_\tau = \frac{1}{p} \left[\frac{p\pi}{c} - \frac{2}{q_{\tau+1}} \right],$$

and hence

$$\omega_{\tau+1} = \frac{1}{q_{\tau+1}} + \frac{1}{p} \left[\frac{p\pi}{c} - \frac{2}{q_{\tau+1}} \right],$$

and therefore

$$\omega_\tau = \frac{\pi}{c} - \left(\frac{2}{p} - 1 \right) \frac{1}{q_\tau},$$

and also

$$\Lambda_{\tau+1} = \frac{\frac{p\pi}{c} - \frac{2}{q_{\tau+1}}}{\delta \left[\frac{p\pi}{c} - \frac{2-p}{q_\tau} \right]},$$

so that

$$\delta\Lambda_\tau = \frac{\psi - 2l_1\beta^\tau}{2 + \psi - (1 + \beta)(1 + Q_1\beta^{\tau-1})},$$

as well as

$$\omega_\tau = \frac{2 + \psi - (1 + \beta)(1 + Q_1\beta^\tau)}{1 - \beta}.$$

Note that delay Λ_τ must be less than one. Rearranging the expression above, this gives

$$q_\tau \leq \frac{2 - (2 - \delta) \frac{p}{1-\delta}}{\psi + 2 - \frac{2p}{1-\delta}} =: q_\Delta^{**} \rightarrow \frac{2 - \sigma}{v + 2 - 2\sigma} =: q^{**}.$$

Also, $(\Lambda_{\tau+2} - \Lambda_{\tau+1}) - (\Lambda_{\tau+1} - \Lambda_\tau)$ is positively proportional to

$$\frac{2\beta - \psi}{1 + \psi - \beta - l_1\beta^{\tau+2}(2 + \beta)},$$

whenever $\Lambda_\tau < 1$. It follows that delay is decreasing in τ for $q < q_\Delta^{**}$ (when $\psi < 2$, $\sigma < \psi$), so that delay is well-defined and positive there. Alternatively, we can appeal to the limiting delay function λ given by (14), noting that $\lambda(q) \geq 1$ over $[q, q^{**}]$, as the function λ given by (14) is decreasing in q over the interval $(q, 1)$ in this case. It is also simple to check that the payoffs of the principal and the agent are positive above q^{**} in

this case.²² To show that the equilibrium outcome must have this structure, note first that we cannot have $\lambda(q) \geq 1$, and hence cannot have delay, for values of q strictly above q^{**} . Finally, we can show that there can be no interval without delay involving values of q less than q^{**} .²³

Conversely, delay is decreasing in τ for $q > q_{\Delta}^*$ (when $\psi > 2$, $\sigma > \psi$). It is not hard to show that the value of λ exceeds 1 on $(q^*, 1]$.²⁴ In fact, delay does not vary continuously at $q = q^*$, i.e. $\lim_{q \downarrow q^*} \lambda(q) > 1$. The uniqueness of this outcome follows from the fact that there cannot be delay for q close to \underline{q} (as (15) is violated at $q = \underline{q}$). Hence, the principal's payoff is given by (2) for beliefs that are low enough, and it then follows by continuity of the principal's payoff that there cannot be delay for $q < q^*$, at which point the function given by (9) dips below zero and delay arises.²⁵

Combined with the observations following the derivation of the sequence ν_{τ} , we obtain that, depending on $\psi \leq 2$, and $\psi \leq \sigma$, four cases can occur. We have thus established the following:

Lemma 17 *Given ψ and σ , there exists $\bar{\Delta}$ such that if $\Delta < \bar{\Delta}$, there exists an equilibrium with*

²²This requires a calculation. First, we can solve for the differential equations giving v and w over the range $[q^{**}, 1]$, where we use as boundary conditions $w(q^{**}) = (q^{**}(\psi + 1) - 2/\sigma)c/r$, and $v(q^{**}) = 0$. It is easy to check that $v''(q^{**}) = 0$ (compare, for instance, with $v''(\underline{q})$ above), so the curvature of v is actually zero at q^{**} . However,

$$v'''(q^{**}) = \frac{\sigma^3(\psi - 1)^5}{(\sigma - 2)^2(2 + \sigma - \sigma(\psi + 1))^2} \frac{c}{r},$$

which is strictly positive. Since v admits at most one inflection point over the interval $(q^{**}, 1)$, and it is positive at 1, it follows that it is positive over the entire interval.

²³That is, assume for the sake of contradiction that there is delay on a non-degenerate interval $[q_1, q_2]$ with $q_1 < q^{**}$. Then since $v(q_1) = 0$ and $w(q_1) = ((q_1(\psi + 2) - 2)/\sigma)(c/r)$, we can solve for v and w , which gives $v'(q_1) = 0$, while

$$v''(q_1) = \left(\frac{\sigma - 2 + \sigma q_1(\psi - 1)}{q_1^2(1 - q_1)^2 \sigma^2} \right) \frac{c}{r}.$$

This value is strictly negative because $\psi < 2$ and $\psi > \sigma$ imply $\sigma - 2 + \sigma q_1(\psi - 1) < 0$. This implies that v is strictly decreasing at q_1 , and hence strictly negative over some range above q_1 , a contradiction.

²⁴It is easy to check that the coefficient of $((1 - q)\underline{q})/(q(1 - \underline{q}))$ in (9) is positive given $\psi > 2$ and $\psi < \sigma$, so that, by ignoring this term while solving for the root $v(q) = 0$, we obtain a lower bound on q^* . That is, $q^* \geq \tilde{q} := (\sigma - 2)(\sigma + 1)/[(\psi - 2)\sigma]$. Since $\lambda(q) > 1$ if and only if $q > q^{**}$ (from (15) given that $\psi < 2$ in this case), it suffices now to note that $\tilde{q} \geq q^{**}$.

²⁵This is not enough to imply that there is delay for all beliefs above q^* . To prove that there cannot be a subinterval (q_1, q_2) of $(q^*, 1]$ in which there is no delay, consider such a maximal such interval and note that it would have to be the case that $v(q_1) = 0$ and $w(q_1) = (q_1\pi - 2c/p)c/r$, by continuity in q of the players' payoff functions. Solving for the differential equations for v, w in such an interval (q_1, q_2) , one obtains that, at q_1 , $v(q_1) = v'(q_1) = 0$, while $v''(q_1)$ is given in footnote 23. The numerator of v'' is necessarily negative for all $q_1 > \tilde{q}$ (cf. footnote 24), and thus, in particular, for $q > q^*$. This contradicts the fact that v must be nonnegative on the interval (q_1, q_2) .

1. No delay (for any $q \in [\underline{q}, 1]$) if $\psi > \sigma$, $\psi > 2$.
2. Delay if and only if $q > q_\Delta^*$, where $q_\Delta^* \rightarrow q^*$ (cf. Lemma 9.2), if $\psi < \sigma$, $\psi > 2$.
3. No Delay for $q \in [\underline{q}, \hat{q}_\Delta]$, delay for $q \in (\hat{q}_\Delta, q_\Delta^{**}]$, and no delay if $q > q_\Delta^{**}$ if $\psi > \sigma$, $\psi < 2$, where $\hat{q}_\Delta \rightarrow \underline{q}$, and $q_\Delta^{**} \rightarrow q^{**}$.
4. No Delay for $q \in [\underline{q}, \hat{q}_\Delta]$, and delay for all $q > \hat{q}_\Delta$, if $\psi < \sigma$, $\psi < 2$, where again $\hat{q}_\Delta \rightarrow \underline{q}$.

To show that in the final case there cannot be an interval (q_1, q_2) in which there is no delay, we again proceed as in footnote 25.²⁶

A.8.3 Summary: Canonical Markov Equilibrium

We can summarize these results in the following proposition.

Proposition 8 *As $\Delta \rightarrow 0$, the canonical Markov equilibrium approaches a limit whose form depends on the project's parameters as follows:*

- *High Surplus, Patient Projects ($\psi > 2$ and $\psi > \sigma$): The principal makes an offer to the agent at every opportunity, until either achieving a success or until the posterior probability of a good project drops below $\underline{q} = \frac{2c}{p\pi}$. The principal's payoff is positive for all posteriors exceeding \underline{q} .*
- *High Surplus, Impatient Projects ($\psi > 2$ and $\psi < \sigma$): The principal initially continually delays before making each offer to the agent, until the posterior probability drops to a threshold $q^* > \underline{q}$. The principal subsequently makes offers with no delay, until the posterior hits \underline{q} . The principal's expected payoff is zero for $q > q^*$ and positive for $q \in (\underline{q}, q^*)$.*
- *Low Surplus, Patient Projects ($\psi < 2$ and $\psi > \sigma$): The principal initially makes offers at every opportunity, enjoying a positive payoff, until the posterior drops to a threshold $q^{**} > \underline{q}$, at which point the principal introduces delay and commands an expected payoff of zero.*
- *Low Surplus, Impatient Projects ($\psi < 2$ and $\psi < \sigma$): Here the principal delays before making each offer, for every posterior, with a zero expected payoff.*

²⁶We show that, solving the differential equations for v and w , the value of $v''(q_1)$ is negative. Since $v(q_1) = v'(q_1) = 0$, this implies that the payoff of the principal would be strictly negative for values of q slightly above q_1 , a contradiction.

A.8.4 Limit Uniqueness

We complete the argument with a limiting result, with Section A.23 providing the proof:

Lemma 18 *The limiting payoff of any sequence of Markov equilibria, as Δ approaches zero, equals the limiting payoff of the canonical Markov equilibrium.*

A.9 s^W vs. s^S , Agent Offers

What if the agent makes the offers instead of the principal, as in Bergemann and Hege [2]? The agent thus chooses the share s_t , but in doing so must respect the incentive constraint that it be optimal to undertake an experiment whenever experimental funding is advanced.

Consider a posterior q_1 with $\varphi(q_1) < \underline{q} < q_1$, so that there will be only one more experiment before beliefs become sufficiently pessimistic ($\varphi(q) < \underline{q}$) as to halt experimentation. If the agent is expected to work, the incentive constraint is

$$pq_1\pi(1 - s_1) \geq pq_1\pi - c.$$

We can solve this for the threshold

$$pq_1\pi s_1^W = pq_1\pi - c.$$

In equilibrium, the agent will push the principal to her participation constraint by setting share s^* solving

$$pq_1\pi s^* = c.$$

Now suppose the agent faces share s_1 and is expected to shirk, after which we come to the next period with unchanged beliefs, at which point share s^* is offered. Then the agent's incentive constraint is

$$c + \delta pq_1\pi(1 - s^*) \geq pq_1\pi(1 - s_1) + \delta(1 - pq_1) \max\{c, p\varphi(q_1)\pi(1 - s^*)\}.$$

We can use the definition of s^* and rearrange to obtain

$$pq_1\pi s_1^S = (1 - \delta)(pq_1\pi - c) + \delta(1 - pq_1) \max\{c, p\varphi(q_1)\pi(1 - s^*)\}.$$

This suffices to give

$$s_1^S < s_1^W,$$

as long as $(1 - pq_1) \max\{c, p\varphi(q_1)\pi(1 - s^*)\} < pq_1\pi - c$. The fact that $c < pq_1\pi - c$ follows from $q_1 > \underline{q}$. Alternatively, we have

$$(1 - pq_1)p\varphi(q_1)\pi(1 - s^*) = (1 - p)pq_1\pi(1 - s^*) = (1 - p)(pq_1\pi - c) < pq_1\pi - c.$$

We thus have a range of shares $[s_1^S, s_1^W]$ for which the agent will work if expected to do so, and will shirk if expected to do so. Bergemann and Hege [2] choose the value of s_1 in order to ensure the principal has a zero payoff. As long as this value falls in the interior of $[s_1^S, s_1^W]$, as it will for q_1 greater than but close to \underline{q} , there will be multiple equilibria. This in turn feeds into additional opportunities for multiplicity as we work backward from the final period.²⁷

To establish the possibility of $s^S > s^W$, notice that in general the incentive constraint when the agent is expected to work is

$$pq\pi(1-s) + \delta(1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)) = c + \delta W(\mathbf{1}_{\varphi(q)}, q),$$

and the constraint when expected to shirk is

$$c + \delta W(\mathbf{1}_q, q) = pq\pi(1-s) + \delta(1-pq)W(\mathbf{1}_q, \varphi(q)).$$

We can solve for

$$\begin{aligned} pq\pi s^W &= pq\pi - c - \delta W(\mathbf{1}_{\varphi(q)}, q) + \delta(1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)), \\ pq\pi s^S &= pq\pi - c - \delta W(\mathbf{1}_q, q) + \delta(1-pq)W(\mathbf{1}_q, \varphi(q)). \end{aligned}$$

Hence, it suffices to show

$$(1-pq)W(\mathbf{1}_q, \varphi(q)) - W(\mathbf{1}_q, q) > (1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)) - W(\mathbf{1}_{\varphi(q)}, q),$$

or

$$W(\mathbf{1}_q, q) - (1-pq)W(\mathbf{1}_q, \varphi(q)) < W(\mathbf{1}_{\varphi(q)}, q) - (1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)).$$

Now let $\psi > 2$ and $\psi > \sigma > 2 - 2p$, with $p > r$. (It is straightforward to find values of p, c, r , and π satisfying these inequalities.) Then we have a high-surplus, impatient project when the principal makes offers, and a high-return, high-discount project when the agent makes offers (cf. Section 4.4). In either case, the resulting Markov equilibrium features no delay. We can then use the agent's incentive constraint to rewrite the previous inequality as

$$c + \delta \frac{q}{\varphi(q)} W(\mathbf{1}_{\varphi(q)}, \varphi(q)) - (1-pq)W(\mathbf{1}_q, \varphi(q)) < W(\mathbf{1}_{\varphi(q)}, q) - (1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)).$$

We potentially underestimate the second term on the right side, and hence obtain a sufficient inequality for $s^S > s^W$, with the following:

$$c + \delta \frac{q}{\varphi(q)} W(\mathbf{1}_{\varphi(q)}, \varphi(q)) - (1-pq)[c + \delta W(\mathbf{1}_{\varphi(q)}, \varphi(q))] < W(\mathbf{1}_{\varphi(q)}, q) - (1-pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)).$$

²⁷We have not taken account here of the opportunity Bergemann and Hege [2] allow for the principal to advance partial funding, but this will not eliminate this phenomenon.

Now writing $W(\mathbf{1}_{\varphi(q)}, \varphi(q)) = \tilde{W}$, we note that c is added and subtracted on the left, and then rearrange, to obtain

$$pq c + \delta(1 - pq)\tilde{W} < (1 - \delta)\frac{q}{\varphi(q)}\tilde{W} - (1 - pq)\tilde{W},$$

and hence

$$\frac{pq c}{1 - \delta} < \left[\frac{q}{\varphi(q)} - (1 - pq) \right] \tilde{W} = \frac{q}{\varphi(q)} [1 - (1 - p)] \tilde{W},$$

or finally

$$\frac{\varphi(q)c}{1 - \delta} < \tilde{W}. \quad (41)$$

The proof of Lemma 10 in Section A.15 shows that if $p > r$ and $2c/\pi < r$ (so that $pq_1 > r$; notice that we can ensure these inequalities while preserving our previously maintained assumptions by taking π to be sufficiently large), then there will be a nonempty interval of values $[q^\dagger, 1)$ for which this inequality holds, given that (i) the principal makes offers, (ii) continuation play features no delay, and (iii) the principal makes offer s^W at every opportunity. We can then find a value $q \in [q^\dagger, 1)$ such that (41) holds at q when the principal makes offers. But the agent's payoff must if anything be larger when the agent rather than the principal makes offers, since the agent will lower s at every opportunity to reduce the principal's payoff to zero, and hence (41) must hold when the agent makes offers, giving the result.²⁸

A.10 Proof of Lemma 4

Fix a prior \bar{q} and waiting time Δ , both hereafter to be omitted from the notation. Notice first that if $\mathbb{E}[q^P] < \frac{c}{p\pi}$, every continuation equilibrium outcome must give the principal a negative payoff. In particular, the total expected surplus under the first-best policy is negative (cf. Section 2.3), and hence so must be the principal's payoff.

Suppose the result is false. Then there exists sequences of integers $\{k(n)\}_{n=1}^\infty$, times $t(n)$, strategy profiles $\sigma(n)$, and histories $h_{t(n)}^P(n)$ such that

1. each $\sigma(n)$ is an equilibrium,
2. each history $h_{t(n)}^P(n)$ arises with positive probability under equilibrium σ_n , has length $t(n)$, and features $k(n)$ offers (and failures),
3. $\lim_{n \rightarrow \infty} k_n = \infty$,

²⁸Section 4.4 noted that shifting the bargaining power from the principal to the agent need not make the agent better off, but it must do so when the parameters are such that we have a high-surplus, impatient project when the principal makes offers, and a high-return, high-discount project when the agent makes offers, so that there is delay in neither circumstance.

$$4. \mathbb{E} \left[q^P | h_{t(n)}^P(n), \sigma(n) \right] > \frac{c}{p\pi},$$

where $\mathbb{E} \left[q^P | h_{t(n)}^P(n), \sigma(n) \right]$ denotes the principal's expectation of the agent's belief at history $h_{t(n)}^P(n)$ under equilibrium $\sigma(n)$. This in turn implies that by taking a subsequence and renumbering, we can construct sequences that preserve these properties as well as choose sequences of integers $\kappa'(n)$ and $\kappa''(n)$ such that $\kappa''(n) - \kappa'(n) > n$, $\frac{1}{n} > \mathbb{E} \left[q^P | h_{t(n)|\kappa'(n)}^P, \sigma(n) \right] - \mathbb{E} \left[q^P | h_{t(n)|\kappa''(n)}^P, \sigma(n) \right]$, and history $h^P(t_n)$ features n offers between periods $\kappa'(n)$ and $\kappa''(n)$. Hence, we must be able to find equilibria with the property that over arbitrarily long sequences of failures, beliefs change arbitrarily little. However, this ensures that for sufficiently large n , the principal's payoff upon reaching history $h_{t(n)|\kappa'_n}^P$ must be negative. In particular, the probability that any single subsequent offer made by the principal in this sequence induces effort is converging uniformly to zero (since otherwise $\frac{1}{n} > \mathbb{E} \left[q^P | h_{t(n)|\kappa'(n)}^P, \sigma(n) \right] - \mathbb{E} \left[q^P | h_{t(n)|\kappa''(n)}^P, \sigma(n) \right]$ is impossible), while every such offer incurs a cost of c . A negative expected payoff for the principal at history $h_{t(n)|\kappa'_n}^P$ is a contradiction.

A.11 Proof of Lemma 5

The game begins with $q^A = \bar{q}$ and with q^P placing unitary probability on \bar{q} . As long as the agent chooses pure actions, the public belief q^P continues to attach unitary probability to a single belief. The first time the agent mixes between working and shirking, the public belief subsequently puts positive probability on two posteriors, say q and $\varphi(q)$.

Our method of proof is to argue that given any two such beliefs, if the agent characterized by posterior $\varphi(q)$ has a weak incentive to work, then the agent characterized by q has a strict incentive to work. This ensures that q^P never attaches positive probability to more than two beliefs. In particular, once two such beliefs have arisen, in the subsequent period either both are revised downward, with q^P then again attaching positive probability to two beliefs (one the Bayesian update of the other); or the smaller belief is subject to no revision, in which case q^P attaches positive probability to at most (the same) two beliefs.

We work backward from the end of the game. Hence, let us renumber the sequence of offers as s_1, s_2, \dots , where s_1 is the last offer made, s_2 the penultimate offer, and so on. We let δ_τ be the discount factor relevant when s_τ is offered. The magnitude of δ_τ will depend on the time that elapses between offer s_τ and offer $s_{\tau-1}$.

A.11.1 The Final Period

Suppose we are in the final period, with share s_1 offered. Then it is obvious that if an agent with belief q_1 finds it optimal to work, so will any agent with belief $q_2 > q_1$.

A.11.2 The Penultimate Period

Now suppose we are in the penultimate period, facing share s_2 , and consider agents with beliefs q_1 and q_2 , with $q_0 = \varphi(q_1)$ and $q_1 = \varphi(q_2)$. Hence, q_0 is the update of q_1 and q_1 is the update of q_2 .

We argue that it is impossible that q_1 would prefer to work and q_2 to shirk, i.e., that it is impossible that

$$\begin{aligned} pq_1\pi(1-s_2) + \delta_2(1-pq_1)\max\{c, pq_0\pi(1-s_1)\} &\geq c + \delta_2\max\{c, pq_1\pi(1-s_1)\}, \\ pq_2\pi(1-s_2) + \delta_2(1-pq_2)\max\{c, pq_1\pi(1-s_1)\} &\leq c + \delta_2\max\{c, pq_2\pi(1-s_1)\}. \end{aligned}$$

The value of s_1 may be random. However, we will argue that for no value of s_1 can these constraints be satisfied. If so, then they cannot be satisfied on average. This suffices to establish the result.

We consider four cases:

Case 1: $c \geq pq_2\pi(1-s_1)$. The incentive constraints are

$$\begin{aligned} pq_1\pi(1-s_2) + \delta_2(1-pq_1)c &\geq c + \delta_2c, \\ pq_2\pi(1-s_2) + \delta_2(1-pq_2)c &\leq c + \delta_2c. \end{aligned} \tag{42}$$

This requires

$$pq_1\pi(1-s_2) + \delta_2(1-pq_1)c \geq pq_2\pi(1-s_2) + \delta_2(1-pq_2)c,$$

or

$$pq_1[\pi(1-s_2) - \delta_2c] \geq pq_2[\pi(1-s_2) - \delta_2c].$$

Since $\pi(1-s_2) - \delta_2c > 0$ (from (42)), this is a contradiction.

Case 2: $c \in [pq_1\pi(1-s_1), pq_2\pi(1-s_2)]$. The incentive constraints are

$$\begin{aligned} pq_1\pi(1-s_2) + \delta_2(1-pq_1)c &\geq c + \delta_2c, \\ pq_2\pi(1-s_2) + \delta_2(1-pq_2)c &\leq c + \delta_2pq_2\pi(1-s_1), \end{aligned}$$

or

$$\begin{aligned} pq_1\pi(1-s_2) &\geq c + \delta_2c - \delta_2(1-pq_1)c, \\ pq_1\pi(1-s_2) &\leq \frac{q_1}{q_2}c + \delta_2pq_1\pi(1-s_1) - \delta_2\frac{q_1}{q_2}(1-pq_2)c. \end{aligned}$$

Hence we need

$$c + \delta_2c - \delta_2(1-pq_1)c \leq \frac{q_1}{q_2}c + \delta_2pq_1\pi(1-s_1) - \delta_2\frac{q_1}{q_2}(1-pq_2)c,$$

or, removing some common terms,

$$c \leq \frac{q_1}{q_2}c + \delta_2 pq_1 \pi(1 - s_1) - \delta_2 \frac{q_1}{q_2}c.$$

This is

$$q_2c + \delta_2 q_1c \leq q_1c + \delta_2 pq_1 \pi(1 - s_1)q_2.$$

We overestimate the right side by writing

$$q_2c + \delta_2 q_1c \leq q_1c + \delta_2 q_2c,$$

which is

$$(1 - \delta_2)q_2 \leq (1 - \delta_2)q_1,$$

a contradiction.

Case 3: $c \in [pq_0\pi(1 - s_1), pq_1\pi(1 - s_1)]$. The incentive constraints are

$$\begin{aligned} pq_1\pi(1 - s_2) + \delta_2(1 - pq_1)c &\geq c + \delta_2 pq_1\pi(1 - s_1), \\ pq_2\pi(1 - s_2) + \delta_2(1 - pq_2)pq_1\pi(1 - s_1) &\leq c + \delta_2 pq_2\pi(1 - s_1), \end{aligned}$$

or

$$\begin{aligned} pq_1\pi(1 - s_2) &\geq c + \delta_2 pq_1\pi(1 - s_1) - \delta_2(1 - pq_1)c, \\ pq_1\pi(1 - s_2) &\leq \frac{q_1}{q_2}c + \delta_2 pq_1\pi(1 - s_1) - \delta_2 \frac{q_1}{q_2}(1 - pq_2)pq_1\pi(1 - s_1). \end{aligned}$$

Hence, we need

$$c + \delta_2 pq_1\pi(1 - s_1) - \delta_2(1 - pq_1)c \leq \frac{q_1}{q_2}c + \delta_2 pq_1\pi(1 - s_1) - \delta_2 \frac{q_1}{q_2}(1 - pq_2)pq_1\pi(1 - s_1),$$

or, eliminating common terms and multiplying by q_2 ,

$$q_2c - \delta_2(1 - pq_1)cq_2 \leq q_1c - \delta_2 q_1(1 - pq_2)pq_1\pi(1 - s_1).$$

We overestimate the right side by writing this as

$$q_2c - \delta_2(1 - pq_1)cq_2 \leq q_1c - \delta_2 q_1(1 - pq_2)c,$$

which is,

$$q_2 - \delta_2 q_2 < q_1 - \delta_2 q_1,$$

a contradiction.

Case 4: $c \leq pq_0\pi(1 - s_1)$. The incentive constraints are

$$\begin{aligned} pq_1\pi(1 - s_2) + \delta_2(1 - pq_1)pq_0\pi(1 - s_1) &\geq c + \delta_2pq_1\pi(1 - s_1), \\ pq_2\pi(1 - s_2) + \delta_2(1 - pq_2)pq_1\pi(1 - s_1) &\leq c + \delta_2pq_2\pi(1 - s_1), \end{aligned}$$

or

$$\begin{aligned} pq_1\pi(1 - s_2) &\geq c + \delta_2pq_1\pi(1 - s_1) - \delta_2(1 - pq_1)pq_0\pi(1 - s_1), \\ pq_1\pi(1 - s_2) &\leq \frac{q_1}{q_2}c + \delta_2pq_1\pi(1 - s_1) - \delta_2\frac{q_1}{q_2}(1 - pq_2)pq_1\pi(1 - s_1). \end{aligned}$$

Hence, we need

$$c + \delta_2pq_1\pi(1 - s_1) - \delta_2(1 - pq_1)pq_0\pi(1 - s_1) \leq \frac{q_1}{q_2}c + \delta_2pq_1\pi(1 - s_1) - \delta_2\frac{q_1}{q_2}(1 - pq_2)pq_1\pi(1 - s_1),$$

or

$$q_2c - q_2\delta_2(1 - pq_1)pq_0\pi(1 - s_1) \leq q_1c - \delta_2q_1(1 - pq_2)pq_1\pi(1 - s_1),$$

and hence

$$q_2c + \delta_2q_1(1 - pq_2)pq_1\pi(1 - s_1) \leq q_1c + q_2\delta_2(1 - pq_1)pq_0\pi(1 - s_1).$$

Since $q_2c > q_1c$, it suffices for a contradiction to show

$$\delta_2q_1(1 - pq_2)pq_1\pi(1 - s_1) \geq q_2\delta_2(1 - pq_1)pq_0\pi(1 - s_1),$$

or

$$q_1(1 - pq_2)q_1 \leq q_2(1 - pq_1)q_0.$$

Using Bayes' rule, this is

$$(1 - p)q_2q_1 \leq (1 - p)q_1q_2,$$

which is obvious, and hence yields the contradiction.

A.11.3 The Induction Step

Now we examine the induction step. We suppose that we are facing offer s_τ . The induction hypothesis is that there is no future offer $\{s_1, \dots, s_{\tau-1}\}$ that induces an agent to work while a more optimistic agent shirks.

Suppose that when making offer s_τ , q^P attaches positive probability to two beliefs q_1 and q_2 , with q_1 being the belief reached from q_2 via updating in the event of a failure. We claim that it is impossible that q_1 works while q_2 shirks, i.e., that it is impossible that

$$pq_1\pi(1 - s_\tau) + \delta(1 - pq_1)W^0 \geq c + \delta W^1 \tag{43}$$

$$pq_2\pi(1 - s_\tau) + \delta(1 - pq_2)W^1 \leq c + \delta W^2, \tag{44}$$

where q_0 is the belief reached from q_1 via Bayesian updating after a failure, W^0 is the continuation value of an agent with posterior q^0 who faces the subsequent sequence of offers, and W^1 and W^1 are analogous for priors q_1 and q_2 . The sequence of shares offered to the agent may be random, in which case these are the appropriate expected values.

Rearranging, we need to show the impossibility of

$$\begin{aligned} pq_1\pi(1-s_\tau) &\geq c + \delta_\tau W^1 - \delta_\tau(1-pq_1)W^0 \\ pq_1\pi(1-s_\tau) &\leq \frac{q_1}{q_2}c + \delta_\tau \frac{q_1}{q_2}W^2 - \delta_\tau \frac{q_1}{q_2}(1-pq_2)W^1. \end{aligned}$$

Given that we have placed no restrictions on s_τ , demonstrating this impossibility is equivalent to showing (now phrasing things positively rather than seeking a contradiction)

$$c + \delta W^1 - \delta_\tau(1-pq_1)W^0 > \frac{q_1}{q_2}c + \delta_\tau \frac{q_1}{q_2}W^2 - \delta_\tau \frac{q_1}{q_2}(1-pq_2)W^1.$$

We can rewrite, using the updating rules, as

$$c - \frac{q_1}{q_2}c + \delta_\tau \left(W^1 - \frac{q_1}{q_2}W^2 \right) > \delta_\tau(1-p) \frac{q_1}{q_0} \left(W^0 - \frac{q_0}{q_1}W^1 \right).$$

The terms W^0 , W^1 and W^2 are sums of equal numbers of terms, one for each offer remaining. Any given offer is common to the three sums, but the actions invoked by a given offer may differ across the sums. By the induction hypothesis, the possible action configurations that might appear in any particular period of the continuation play generating W^0 , W^1 and W^2 , respectively, are sss , ssw , sww , and www .

Now let us suppose that sss occurs in response to some future offer, at which point continuation paths W^0 , W^1 , and W^2 have hit posterior beliefs \tilde{q}_0 , \tilde{q}_1 and \tilde{q}_2 , respectively, and that all previous periods have featured www . We argue that the contribution of this future period to the inequality

$$\delta_\tau \left(W^1 - \frac{q_1}{q_2}W^2 \right) \geq \delta_\tau(1-p) \frac{q_1}{q_0} \left(W^0 - \frac{q_0}{q_1}W^1 \right)$$

satisfies this inequality. In particular, this contribution is (after eliminating some common terms)

$$(1-pq_1)(1-p\tilde{q}_1) - \frac{q_1}{q_2}(1-pq_1)(1-pq_2) \geq (1-p) \frac{q_1}{q_0}(1-p\tilde{q}_1)(1-p\tilde{q}_0) - (1-p)(1-pq_1)(1-p\tilde{q}_1).$$

Using the updating rule, this is

$$(1-pq_1)(1-p\tilde{q}_1) - (1-p)(1-pq_1) \geq (1-pq_1)(1-p\tilde{q}_1)(1-p\tilde{q}_0) - (1-p)(1-pq_1)(1-p\tilde{q}_1).$$

Deleting the common $(1 - pq_1)$ gives

$$(1 - p\tilde{q}_1) - (1 - p) \geq (1 - p\tilde{q}_1)(1 - p\tilde{q}_0) - (1 - p)(1 - p\tilde{q}_1).$$

Collecting terms, we have

$$1 - p\tilde{q}_1 \geq (1 - p\tilde{q}_1)(1 - p\tilde{q}_0) + (1 - p)p\tilde{q}_1,$$

or

$$p\tilde{q}_0(1 - p\tilde{q}_1) \geq p(1 - p)\tilde{q}_1,$$

or

$$\tilde{q}_0(1 - p\tilde{q}_1) \geq (1 - p)\tilde{q}_1,$$

which holds as an equality, as a restatement of the updating rule.

This means that we can effectively remove from consideration any action profile sss that appears before the first sww or ssw , replacing the play of sss by inaction and an appropriately reduced discount factor to capture the passage of time from the offer preceding the (removed) instance of sss to the following offer. We thus need only consider paths of play featuring a succession of periods of www followed by sww , or a succession of periods of www followed by a period of ssw (and then some continuation).

Consider the former. We now note that

$$W^1 \geq \frac{q_1}{q_2}W^2,$$

which follows from the fact that $\frac{q_1}{q_2} < 1$ and the pessimistic agent can always mimic the actions of the more optimistic agent. Hence, it suffices to show that

$$c - \frac{q_1}{q_2}c > \delta_\tau(1 - p)\frac{q_1}{q_0} \left(W^0 - \frac{q_0}{q_1}W^1 \right).$$

We focus on the worst case by taking $\delta_\tau = 1$, in which case it suffices to show a weak inequality in the preceding relationship. We now claim

$$W^0 - \frac{q_0}{q_1}W^1 \leq c - \frac{q_0}{q_1}c. \tag{45}$$

Let us first consider the implications of this inequality. It implies that it suffices to establish

$$c - \frac{q_1}{q_2}c \geq (1 - p)\frac{q_1}{q_0} \left(c - \frac{q_0}{q_1}c \right).$$

Successive simplifications give

$$\begin{aligned}
1 - \frac{q_1}{q_2} &\geq (1 - p) \frac{q_1}{q_0} - (1 - p) \\
2 - \frac{1 - p}{1 - pq_2} &\geq (1 - pq_1) + p \\
2 - 2pq_2 - 1 + p &\geq 1 - pq_1 - pq_2 + p^2 q_1 q_2 + p - p^2 q_2 \\
-pq_2 &\geq p^2 q_1 q_2 - p^2 q_2 - pq_1 \\
q_1 + pq_2 &\geq q_2 + pq_1 q_2 \\
1 + p \frac{q_2}{q_1} &\geq \frac{q_2}{q_1} + pq_2 \\
1 - pq_2 &\geq (1 - p) \frac{q_2}{q_1} = 1 - pq_2,
\end{aligned}$$

which obviously holds.

So, we need to establish (45). By assumption, the paths inducing W^0 and W^1 both feature effort in response to offers $s_{\tau-1}$ through $s_{\tilde{\tau}}$ for some $\tilde{\tau}$. In responding to each of these offers, the payoff under W^1 is precisely $\frac{q_1}{q_0}$ that of W^0 , and hence these periods contribute nothing to the difference $W^0 - \frac{q_0}{q_1} W^1$. In the next period, W^0 shirks while W^1 exerts effort. Letting \hat{W}^0 and \hat{W}^1 denote the continuation values beginning in period $\tilde{\tau} - 1$, we can write (using the incentive constraint for the second relationship)

$$\begin{aligned}
\hat{W}^0 &= \left(\prod_{s=\tau-1}^{\tilde{\tau}} \delta_s \right) \left(\prod_{s=0}^{\tau-\tilde{\tau}} (1 - pq_{-s}) \right) (c + \delta_{\tilde{\tau}-1} \tilde{W}) \\
\hat{W}^1 &\geq \left(\prod_{s=\tau-1}^{\tilde{\tau}} \delta_s \right) \left(\prod_{s=0}^{\tau-\tilde{\tau}} (1 - pq_{1-s}) \right) \left(c + \delta_{\tilde{\tau}-1} \frac{q_{-(\tau-\tilde{\tau})}}{q_{-(\tau-\tilde{\tau})-1}} \tilde{W} \right),
\end{aligned}$$

where \tilde{W} is the same in both cases. The contribution to the expression $W^0 - \frac{q_0}{q_1} W^1$ given by terms involving \tilde{W} is then, for some constant K ,

$$K \left((1 - pq_{-(\tau-\tilde{\tau})}) - (1 - pq_1) \frac{q_{-(\tau-\tilde{\tau})}}{q_{-(\tau-\tilde{\tau})-1}} \frac{q_0}{q_1} \right) \tilde{W},$$

which, using the rules for belief updating, equals zero. Hence, we have contributions to the difference only from terms involving c . If we want to maximize the contribution to the difference from these terms, we should examine cases in which this shirking happens in the first period. This gives us an upper bound that matches (45).

The remaining possibility to be considered is that we have a succession of periods of www followed by ssw . Here, a direct calculation shows that (43)–(44) cannot both hold. Consider the agent with belief q_2 . Our putative equilibrium behavior calls for this agent to shirk in period τ , and then work through period $\tilde{\tau}$, and shirk in period $\tilde{\tau} - 1$. If (44)

is to hold, it must hold when we consider the alternative course of action in which player q_2 works for in periods τ through $\tilde{\tau}$, and then shirks in period $\tilde{\tau} - 1$. Hence, if (44) is to hold, we must have

$$\begin{aligned}
& c + \delta_\tau p q_2 \pi (1 - s_\tau) + \delta_\tau \delta_{\tau-1} (1 - p) p q_2 \pi (1 - s_{\tau-1}) + \\
& \quad \cdots + \prod_{s=0}^{\tau-\tilde{\tau}} \delta_{\tau-s} (1 - p)^{\tau-\tilde{\tau}} p q_2 \pi (1 - s_{\tilde{\tau}}) \\
\geq & p q_2 \pi (1 - s_\tau) + \delta_{\tau-1} (1 - p) p q_2 \pi (1 - s_{\tau-1}) + \\
& \quad \cdots + \prod_{s=0}^{\tau-\tilde{\tau}-1} \delta_{\tau-s} (1 - p)^{\tau-\tilde{\tau}-1} + \prod_{s=0}^{\tau-\tilde{\tau}} \delta_{\tau-s} (1 - p)^{\tau-\tilde{\tau}} c.
\end{aligned}$$

Notice that no period after $\tau' - 1$ enters these payoff calculations. The two paths under consideration yield identical payoffs in later periods, and hence these periods can be neglected.

Similarly, consider the player characterized by belief q_1 . In the putative equilibrium, this player works for some number k of periods and then shirks. If (43) is to hold, it must hold for the alternative continuation path in which player q_1 first shirks and then works for k periods. Again, these two paths lead to identical payoffs in periods beyond the first $k + 1$. The requirement that (43) hold is then

$$\begin{aligned}
& c + \delta_\tau p q_1 \pi (1 - s_\tau) + \delta_\tau \delta_{\tau-1} (1 - p) p q_1 \pi (1 - s_{\tau-1}) + \\
& \quad \cdots + \prod_{s=0}^{\tau-\tau'} \delta_{\tau-s} (1 - p)^{\tau-\tau'} p q_1 \pi (1 - s_{\tau'}) \\
\geq & p q_1 \pi (1 - s_\tau) + \delta_{\tau-1} (1 - p) p q_1 \pi (1 - s_{\tau-1}) + \\
& \quad \cdots + \prod_{s=0}^{\tau-\tau'-1} \delta_{\tau-s} (1 - p)^{\tau-\tau'-1} + \prod_{s=0}^{\tau-\tau'} \delta_{\tau-s} (1 - p)^{\tau-\tau'} c.
\end{aligned}$$

We can rewrite these as

$$\begin{aligned}
c & \geq q_2 H, \\
c & \leq q_1 H,
\end{aligned}$$

for some $H > 0$. Since $q_2 > q_1$, this is a contradiction.

A.12 Proof of Lemma 6

We invoke a simple induction argument. To do this, let τ index the number of offers still to be made along the equilibrium path, i.e., the number of failures that will be endured until play ceases. Suppose we have reached the last offer s_1 (and hence $\tau = 1$) of the game, and that $(q^P, q^A) = (\mathbf{1}_{q_1}, q_1)$. In equilibrium, the agent's value is then

$$W(\mathbf{1}_{q_1}, q_1) = (1 - s_1) q_1 p \pi \geq c,$$

where the inequality is the incentive constraint that the agent want to work, devoid of a continuation value in this case because there is no continuation. Now observe that if the agent holds the private belief $\tilde{q} > q_1$, then again the agent will be asked to work one period. Hence,

$$\begin{aligned}
W(\mathbf{1}_{q_1}, \tilde{q}) &= p\tilde{q}\pi(1 - s_1) \\
&= \frac{\tilde{q}}{q_1}pq_1\pi(1 - s_1) \\
&= \frac{\tilde{q}}{q_1}W(\mathbf{1}_{q_1}, q_1) \\
&> c,
\end{aligned}$$

where the final inequality provides the (strict) incentive constraint, ensuring that the agent will indeed work.

Now suppose we have reached a history in which, in equilibrium, there are τ periods to go, with beliefs $(\mathbf{1}_{q_\tau}, \tilde{q})$ for $q_\tau < \tilde{q}$, and suppose that (32) holds for all periods $\tilde{\tau} < \tau$. Then we have

$$\begin{aligned}
W(\mathbf{1}_{q_\tau}, \tilde{q}) &= p\tilde{q}\pi(1 - s_\tau) + \delta_\tau(1 - p\tilde{q})W(\mathbf{1}_{q_{\tau-1}}, \varphi(\tilde{q})) \\
&= \frac{\tilde{q}}{q_\tau} \left[pq_\tau\pi(1 - s_\tau) + \delta_\tau(1 - p\tilde{q}) \frac{q_\tau}{\tilde{q}} \frac{\varphi(\tilde{q})}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \right] \\
&= \frac{\tilde{q}}{q_\tau} \left[pq_\tau\pi(1 - s_\tau) + \delta_\tau(1 - p\tilde{q}) \frac{1 - q_\tau p}{1 - \tilde{q}p} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \right] \\
&= \frac{\tilde{q}}{q_\tau} [pq_\tau\pi(1 - s_\tau) + \delta_\tau(1 - pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})] \\
&= \frac{\tilde{q}}{q_\tau} W(\mathbf{1}_{q_\tau}, q_\tau),
\end{aligned}$$

where the second equality uses the induction hypothesis, the third invokes the definition of the updating rule φ , the fourth rearranges terms, and the final equality uses the definition of W .

This argument assumes that, given the equilibrium hypothesis that the agent will work in every period, an agent who arrives in period τ with posterior $\tilde{q} > q_\tau$ will find it optimal to work. This follows from Lemma 5.

A.13 Proof of Lemma 7

We offer an induction argument. Let q_τ identify the belief when there are τ periods to go, so that $q_{\tau-1}$ is derived from q_τ via Bayesian updating (given a failure).

In the last period, we have

$$W(\mathbf{1}_{q_1}, q_0) = c = c + \delta_1 W(\mathbf{1}_{q_0}, q_0),$$

since an agent will shirk in the last period if too pessimistic to work.

In the final two periods, we have the equilibrium payoffs

$$\begin{aligned}
W(\mathbf{1}_{q_1}, q_1) &= pq_1\pi(1 - s_1) \\
&= c, \\
\\
W(\mathbf{1}_{q_2}, q_2) &= pq_2\pi(1 - s_2) + \delta_2(1 - pq_2)W(\mathbf{1}_{q_1}, q_1) \\
&= pq_2\pi(1 - s_2) + \delta_2(1 - pq_2)pq_1\pi(1 - s_1) \\
&= pq_2\pi(1 - s_2) + \delta_2(1 - pq_2)c \\
&= c + \delta_2pq_2\pi(1 - s_1),
\end{aligned}$$

where the final equality in each case is the incentive constraint. We then have

$$W(\mathbf{1}_{q_2}, q_1) = \max \begin{cases} c + \delta_2pq_1\pi(1 - s_1) = c + \delta_2c \\ pq_1\pi(1 - s_2) + \delta_2(1 - pq_1)c, \end{cases}$$

where the first line is the value if the agent shirks in the current period, and the next line is the value if the agent waits until the final period to shirk. (Never shirking is clearly suboptimal, as is shirking in both periods.) We will have established the result (for the case of the final two periods) if we show that the first of these is the larger, or

$$c + \delta_2c \geq pq_1\pi(1 - s_2) + \delta_2(1 - pq_1)c.$$

We can eliminate a term δ_2c from both sides to obtain the first inequality in the following and then use the incentive constraint for $W(\mathbf{1}_{q_2}, q_2)$ for the latter:

$$\begin{aligned}
c + \delta_2pq_1c &\geq \frac{q_1}{q_2}pq_2\pi(1 - s_2) \\
&= \frac{q_1}{q_2} [c + \delta_2pq_2\pi(1 - s_1) - \delta_2(1 - pq_2)c].
\end{aligned}$$

This rearranges to

$$c + \delta_2pq_1c + \delta_2(1 - pq_2)\frac{q_1}{q_2}c \geq \frac{q_1}{q_2}c + \delta_2pq_1\pi(1 - s_1),$$

or, noting that the δ_2pq_1c terms on the left cancel and using the last-period incentive constraint,

$$c + \delta_2\frac{q_1}{q_2}c \geq \frac{q_1}{q_2}c + \delta_2c,$$

which reduces to

$$1 \geq \frac{q_1}{q_2},$$

providing the result.

Now we consider an arbitrary period τ , with the current belief being q_τ , and with a failure giving rise to the updated belief $q_{\tau-1}$, and a subsequent failure to the belief $q_{\tau-2}$. We have

$$\begin{aligned} W(\mathbf{1}_{q_\tau}, q_\tau) &= pq_\tau\pi(1-s_\tau) + \delta_\tau(1-pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &= c + \delta_\tau \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \end{aligned}$$

Then we have

$$W(\mathbf{1}_{q_\tau}, q_{\tau-1}) = \max \left\{ \begin{aligned} &c + \delta_\tau W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) = c + \delta_\tau c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2}) \\ &pq_{\tau-1}\pi(1-s_\tau) + \delta_\tau(1-pq_{\tau-1})[c + \delta_{\tau-1}W(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2})], \end{aligned} \right.$$

where the first line is the value if the agent shirks in the current period (with the second equality in this line using the Markov hypothesis to substitute for $W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})$), and the second line is the value if the agent does not shirk in this period. In this case, we use the induction hypothesis, ensuring that the agent will shirk in the next period, allowing us to write $W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-2}) = c + \delta_{\tau-1}W(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2})$.

We show that the former is larger, or (writing $W_{\tau-2}$ for $W(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2})$)

$$c + \delta_\tau c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2} \geq pq_{\tau-1}\pi(1-s_\tau) + \delta_\tau(1-pq_{\tau-1})c + \delta_\tau \delta_{\tau-1}(1-pq_{\tau-1})W_{\tau-2}.$$

We remove a term $\delta_\tau c$ from each side and use the incentive constraint for $W(\mathbf{1}_{q_\tau}, q_\tau)$ to write this as

$$c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2} \geq \frac{q_{\tau-1}}{q_\tau} \left[c + \delta_\tau \frac{q_\tau}{q_{\tau-1}} W_{\tau-1} - \delta_\tau(1-pq_\tau)W_{\tau-1} \right] - \delta_\tau pq_{\tau-1}c + \delta_\tau \delta_{\tau-1}(1-pq_{\tau-1})W_{\tau-2}$$

(where $W_{\tau-1} := W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})$). Now substituting for $W_{\tau-1}$, this is

$$\begin{aligned} c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2} &\geq \frac{q_{\tau-1}}{q_\tau} c + \frac{q_{\tau-1}}{q_\tau} \left[\delta_\tau \frac{q_\tau}{q_{\tau-1}} - \delta_\tau(1-pq_\tau) \right] \left(c + \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2} \right) \\ &\quad - \delta_\tau pq_{\tau-1}c + \delta_\tau \delta_{\tau-1}(1-pq_{\tau-1})W_{\tau-2}. \end{aligned}$$

Expanding, this is

$$\begin{aligned} c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-1}} W_{\tau-2} &\geq \frac{q_{\tau-1}}{q_\tau} c + \delta_\tau c - \delta_\tau \frac{q_{\tau-1}}{q_\tau} (1-pq_\tau)c + \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2} \\ &\quad - \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_\tau} \frac{q_{\tau-1}}{q_{\tau-2}} (1-pq_\tau)W_{\tau-2} - \delta_\tau pq_{\tau-1}c + \delta_\tau \delta_{\tau-1}(1-pq_{\tau-1})W_{\tau-2}. \end{aligned}$$

Each side has a term $\delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_{\tau-2}} W_{\tau-2}$ that can be eliminated, and the term $\delta_\tau p q_{\tau-1} c$ is added and subtracted on the right, which can be eliminated, allowing us to move a $\delta_\tau c$ from the right to the left and obtain

$$(1 - \delta_\tau)c \geq \frac{q_{\tau-1}}{q_\tau}c - \delta_\tau \frac{q_{\tau-1}}{q_\tau}c - \delta_\tau \delta_{\tau-1} \frac{q_{\tau-1}}{q_\tau} \frac{q_{\tau-1}}{q_{\tau-2}}(1 - p q_\tau)W_{\tau-2} + \delta_\tau \delta_{\tau-1}(1 - p q_{\tau-1})W_{\tau-2}.$$

This is

$$(1 - \delta_\tau)c \geq (1 - \delta_\tau)c \frac{q_{\tau-1}}{q_\tau} + \delta_\tau \delta_{\tau-1} W_{\tau-2} \left[(1 - p q_{\tau-1}) - \frac{q_{\tau-1}}{q_\tau} \frac{q_{\tau-1}}{q_{\tau-2}}(1 - p q_\tau) \right],$$

which is verified by noting that the term in brackets on the right is zero.

A.14 Proof of Lemma 9

We study the pointwise limit of v_τ . To do so, we make a change of variable to write the principal's payoff v as a function of the current posterior q_t . We have, to the second order,

$$v(q_t) = [q_t p s(q_t) \pi - c] \Delta + (1 - r \Delta)(1 - p q_t \Delta)v(q_{t+\Delta}).$$

The pointwise limit of this function (as $\Delta \rightarrow 0$) is differentiable, so that it must solve

$$0 = (p q \pi - c) - p q (1 - s(q)) \pi - (r + p q)v(q) - p q (1 - q)v'(q) = 0. \quad (46)$$

Similarly, whenever the agent is indifferent between shirking and not, the *on-path* payoff to the agent, $w(q_t)$, must solve, to the second order,

$$\begin{aligned} w(q_t) &= q_t p (1 - s(q_t)) \pi \Delta + (1 - r \Delta)(1 - q_t p \Delta)w(q_{t+\Delta}) \\ &= c \Delta + (1 - r \Delta)(w(q_{t+\Delta}) + x(q_t) \Delta), \end{aligned}$$

where $x(q_t)$ is the marginal gain from $t + \Delta$ onward from not exerting effort at t (recalling that effort is then optimal at all later dates, since the off-the-equilibrium path relative optimism of the agent makes the agent more likely to accept the principal's offer). Using (32), we obtain

$$x(q_t) \Delta = w(q_{t+\Delta}, q_t) - w(q_{t+\Delta}) = \left(\frac{q_t}{q_{t+\Delta}} - 1 \right) w(q_{t+\Delta}),$$

or, in the limit, given the evolution of the belief q_t under full effort,

$$x(q_t) \Delta = p(1 - q_t)w(q_t) \Delta.$$

Hence, the marginal gain is given, in the limit, by

$$x(q_t) = \int_t^T e^{-\int_t^u r d\tau} (-\dot{q}_u) p(1 - s(q_u)) \pi du. \text{ }^{29}$$

Inserting $x(q_t)\Delta$ into the agent's incentive constraint and taking limits, the agent's payoff satisfies

$$\begin{aligned} 0 &= qp\pi(1 - s(q)) - pq(1 - q)w'(q) - (r + qp)w(q) \\ &= c - rw(q) - pq(1 - q)w'(q) + p(1 - q)w(q). \end{aligned} \quad (47)$$

We now solve for the equilibrium payoffs for this case. Equation (47) reduces to

$$q_t p \pi (1 - s_t) = (r + q_t p) w_t - \dot{w}_t \quad \text{and} \quad r w_t - \dot{w}_t - c = p(1 - q_t) w_t.$$

The second of these equations can be rewritten as

$$r w(q) + pq(1 - q)w'(q) - c = p(1 - q)w(q),$$

where w' is the derivative of w . The solution to this differential equation is

$$w(q) = \frac{pq - r}{p - r} \frac{c}{r} + A(1 - q)^{r/p} q^{1-r/p}, \quad (48)$$

for some constant A . Let $\gamma(q) = pq\pi(1 - s)$ (where, with an abuse of notation, s is a function of q) so the first equation writes

$$\gamma(q) = (r + pq)w(q) + pq(1 - q)w'(q) = \frac{p^2 q - r^2}{p - r} \frac{c}{r} + Ap(1 - q)^{r/p} q^{1-r/p}, \quad (49)$$

giving us the share s . Finally, using the previous equation to eliminate s , equation (46) simplifies to

$$0 = pq\pi - c - \gamma(q) - (r + pq)v(q) - pq(1 - q)v'(q).$$

The solution to this differential equation is given by

$$v(q) = \frac{pq\pi}{p + r} + \frac{2r^2 - p^2 + pr(1 - 2q)}{r(p^2 - r^2)} c + (B(1 - q) - A) \left(\frac{1 - q}{q} \right)^{r/p}, \quad (50)$$

for some constant B . Note that the function $v(q)$ given by (50) yields

$$v(1) = \frac{\psi - \sigma c}{\sigma + 1 r}, \quad (51)$$

²⁹To understand the expression for $x(q_t)$, note that, at any later time u , the agent gets his share $(1 - s(q_u))$ of π with a probability that is increased by a factor $-\dot{q}(u)$, relative to what it would have been had he not deviated. Of course, even if the project is good, it succeeds only at a rate p , and this additional profit must be discounted.

which is positive if and only if $\psi > \sigma$.³⁰

We have thus solved for the payoffs to both agents (given by (48) and (50)), as well as for the share s (given by 49)), over any interval of time featuring no delay. Note that the function v has at most one inflection point in the unit interval, given, if any, by

$$\frac{(A - B)(p + r)}{2pA - (p + r)B},$$

and so it has at most three zeroes. Note also that, if the interval without delay includes \underline{q} , we can solve for the constants of integration A and B using $v(\underline{q}) = w(\underline{q}) = 0$, namely

$$A = \frac{(\sigma \underline{q} - 1) \left(\frac{\underline{q}}{1 - \underline{q}} \right)^{\frac{1}{\sigma}} c}{\underline{q}(1 - \sigma) r} \quad \text{and} \quad B = \frac{[(\psi + 2)^2(1 + \sigma) - 8\sigma(\psi + 1)]\psi^{-1-1/\sigma} c}{4(1 - \sigma^2) r}.$$

Plugging back into the value for v , we obtain that

$$v'(\underline{q}) = 0, \quad v''(\underline{q}) = \frac{(\psi + 2)^3(\psi - 2) c}{4\sigma\psi^2 r}. \quad (53)$$

From (53), v is positive or negative for q close to \underline{q} according to whether v is convex or concave at \underline{q} ; it is positive if

$$\psi > 2,$$

and negative if $\psi < 2$. From (10), $v(1)$ is positive (and hence we can induce full effort and avoid delay for high posteriors) if $\psi > \sigma$ and negative if $\psi < \sigma$. Hence, when $\psi > 2$ and $\psi < \sigma$, v must admit a root $q^* \in (\underline{q}, 1)$. Finally, differentiating (50), we have that $v''(\underline{q}) \neq 0$.

³⁰We can obtain this result directly from (40). As $\tau \rightarrow \infty$, the initial conditions then become insignificant and the principal's payoff approaches

$$\frac{\psi + 1}{1 - (1 - p)\delta}(\psi + 1) - \frac{1}{1 - \delta}.$$

This payoff exceeds zero if

$$\frac{1 - \delta}{1 - (1 - p)\delta}(\psi + 1) > 1. \quad (52)$$

Condition (52) is thus necessary for a no-delay Markov equilibrium, with the candidate equilibrium strategies otherwise driving the principal's payoff below zero for high prior values \bar{q} .

Making the substitutions $\delta = 1 - r\Delta$ and then taking the limit as $\Delta \rightarrow 0$, this rearranges to

$$\psi = \frac{p\pi - 2c}{c} > \frac{p}{r} = \sigma.$$

Then a necessary condition for our candidate strategies to be an equilibrium is that $\psi > \sigma$.

A.15 Proof of Lemma 10

Fix q and simplify the notation by writing $s^W(q)$ and $q^S(q)$ simply as s^W and s^S . We can rearrange the constraints defining s^W and s^S to give

$$\begin{aligned}\frac{1}{\delta}[pq\pi(1 - s^W) - c] &= \frac{q}{\varphi(q)}W(\mathbf{1}_{\varphi(q)}, \varphi(q)) - (1 - pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)) \\ \frac{1}{\delta}[pq\pi(1 - s^S) - c] &= W(\mathbf{1}_q, q) - (1 - pq)W(\mathbf{1}_q, \varphi(q)).\end{aligned}$$

The condition that $s^S < s^W$ is equivalent to

$$W(\mathbf{1}_q, q) - (1 - pq)W(\mathbf{1}_q, \varphi(q)) \geq \frac{q}{\varphi(q)}W(\mathbf{1}_{\varphi(q)}, \varphi(q)) - (1 - pq)W(\mathbf{1}_{\varphi(q)}, \varphi(q)).$$

Substituting for $W(\mathbf{1}_q, q)$ from (34), using Lemma 7, and writing $W(\mathbf{1}_{\varphi(q)}, \varphi(q)) := \tilde{W}$, this is

$$c + \delta \frac{q}{\varphi(q)}\tilde{W} - (1 - pq)[c + \delta\tilde{W}] \geq \frac{q}{\varphi(q)}\tilde{W} - (1 - pq)\tilde{W},$$

or, noting that c is added and subtracted on the left and rearranging,

$$pqc + \delta(1 - pq)\tilde{W} \geq (1 - \delta)\frac{q}{\varphi(q)}\tilde{W} - (1 - pq)\tilde{W},$$

and hence

$$\frac{pqc}{1 - \delta} \geq \left[\frac{q}{\varphi(q)} - (1 - pq) \right] \tilde{W} = \frac{q}{\varphi(q)}[1 - (1 - p)]\tilde{W},$$

or finally

$$\frac{\varphi(q)c}{1 - \delta} \geq \tilde{W}. \tag{54}$$

One case is immediate. Suppose we are in the penultimate period. Then $\tilde{W} = c$, and hence (54) becomes:³¹

$$\frac{\varphi(q)}{1 - \delta} \geq 1.$$

For sufficiently small Δ , and hence large δ , this condition will hold. This establishes that $s^S < s^W$ for small values of q .

³¹We could have derived this result directly. In the penultimate period, the relevant incentive constraints are

$$\begin{aligned}pq\pi(1 - s^W) + \delta(1 - pq)c &= c + \delta \frac{q}{\varphi(q)}c \\ pq\pi(1 - s^S) + \delta(1 - pq)[c + \delta c] &= c + \delta \left[c + \delta \frac{q}{\varphi(q)}c \right].\end{aligned}$$

Returning to (54), think of the equilibrium sequence q_1, q_2, q_3, \dots of posteriors, with q_1 being the smallest posterior larger than the termination boundary \underline{q} , and with $q_{\tau-1}$ obtained from q_τ via Bayes' rule. Let us write $W_\tau := W(\mathbf{1}_{q_\tau}, q_\tau)$, and ask when

$$W_\tau \leq \frac{q_\tau c}{1 - \delta},$$

which will suffice for $s^S(q_t) \leq s^W(q_t)$. We have

$$W_\tau = cq_\tau \left[\frac{1}{q_\tau} + \frac{\delta}{q_{\tau-1}} + \frac{\delta^2}{q_{\tau-2}} + \frac{\delta^3}{q_{\tau-3}} + \dots + \frac{\delta^{\tau-1}}{q_1} \right].$$

This allows us to construct the difference equations

$$W_\tau = \delta \frac{q_\tau}{q_{\tau-1}} W_{\tau-1} + c,$$

and

$$\frac{q_\tau c}{1 - \delta} = \frac{q_\tau}{q_{\tau-1}} \left(\frac{q_{\tau-1} c}{1 - \delta} \right).$$

Now, it suffices to show that $s^S < s^W$ for all q to show that the difference equation giving us W_τ lies everywhere below that for $\frac{q_\tau c}{1 - \delta}$. It is useful to divide W_τ by $\frac{q_\tau c}{1 - \delta}$ (defining Ξ_τ to be the ratio) to get

$$\Xi_\tau = (1 - \delta) \left[\frac{1}{q_\tau} + \frac{\delta}{q_{\tau-1}} + \frac{\delta^2}{q_{\tau-2}} + \frac{\delta^3}{q_{\tau-3}} + \dots + \frac{\delta^{\tau-1}}{q_1} \right],$$

We rearrange to get

$$\begin{aligned} \frac{1}{\delta c} [pq\pi(1 - s^W) - c] &= \frac{q}{\varphi(q)} - (1 - pq), \\ \frac{1}{\delta c} [pq\pi(1 - s^S) - c] &= 1 + \delta \frac{q}{\varphi(q)} - (1 - pq)(1 + \delta). \end{aligned}$$

Then $s^S < s^W$ if

$$\frac{q}{\varphi(q)} - (1 - pq) < 1 + \delta \frac{q}{\varphi(q)} - (1 - pq)(1 + \delta).$$

This rearranges to

$$(1 - \delta) \left[\frac{q}{\varphi(q)} - (1 - pq) \right] < pq,$$

or

$$(1 - \delta) \left[\frac{q}{\varphi(q)} - \frac{q}{\varphi(q)}(1 - p) \right] < pq,$$

which is

$$\frac{\varphi(q)}{1 - \delta} \geq 1.$$

which in turn gives us a difference equation

$$\Xi_{\tau+1} = \delta\Xi_{\tau} + (1 - \delta)\frac{1}{q_{\tau+1}} = \delta\Xi_{\tau} + (1 - \delta)\left(1 + \frac{1 - q_1}{q_1}(1 - p)^{\tau}\right),$$

with initial condition $\Xi_1 = \frac{1-\delta}{q_1}$. We have $s_{\tau+1}^S < s_{\tau+1}^W$ if $\Xi_{\tau} < 1$. We know that $\lim_{\tau \rightarrow \infty} \Xi_{\tau} = 1$, and also that if any τ gives $\Xi_{\tau} > 1$, so do all subsequent τ .

Let us simplify the notation by letting $(1 - q_1)/q_1 = Q_1$ and then write

$$\Xi_{\tau+1} = \delta\Xi_{\tau} + (1 - \delta)(1 + Q_1(1 - p)^{\tau}).$$

We can solve this by writing

$$\Xi_1 = \Xi_1$$

$$\Xi_2 = \delta\Xi_1 + (1 - \delta)(1 + Q_1(1 - p))$$

$$\Xi_3 = \delta^2\Xi_1 + (1 - \delta)(\delta + 1 + \delta Q_1(1 - p) + Q_1(1 - p)^2)$$

$$\Xi_4 = \delta^3\Xi_1 + (1 - \delta)(\delta^2 + \delta + 1 + \delta^2 Q_1(1 - p) + \delta Q_1(1 - p)^2 + Q_1(1 - p)^3)$$

⋮

$$\Xi_{\tau} = \delta^{\tau-1}\Xi_1 + (1 - \delta)(1 + \delta + \dots + \delta^{\tau-2} + Q_1((1 - p)^{\tau-1} + \delta(1 - p)^{\tau-2} + \dots + \delta^{\tau-2}(1 - p))).$$

We can simplify this result to

$$\Xi_{\tau} = \begin{cases} \delta^{\tau-1}\Xi_1 + 1 - \delta^{\tau-1} + (1 - \delta)(1 - p)Q_1\delta^{\tau-2}\frac{1 - \left(\frac{1-p}{\delta}\right)^{\tau-1}}{1 - \frac{1-p}{\delta}} & \delta > 1 - p, \\ \delta^{\tau-1}\Xi_1 + 1 - \delta^{\tau-1} + (1 - \delta)(1 - p)Q_1(1 - p)^{\tau-2}\frac{1 - \left(\frac{\delta}{1-p}\right)^{\tau-1}}{1 - \frac{\delta}{1-p}} & \delta < 1 - p. \end{cases}$$

Suppose first that $p < 1 - \delta$ (or, for arbitrarily small Δ , $p < r$). Then we have

$$\begin{aligned} \Xi_{\tau} &= \delta^{\tau-1}\Xi_1 + 1 - \delta^{\tau-1} + (1 - \delta)Q_1(1 - p)^{\tau-1}\frac{1 - \left(\frac{\delta}{1-p}\right)^{\tau-1}}{1 - \frac{\delta}{1-p}} \\ &= \delta^{\tau-1}\Xi_1 + 1 - \delta^{\tau-1} + (1 - \delta)Q_1(1 - p)\frac{(1 - p)^{\tau-1} - \delta^{\tau-1}}{1 - p - \delta}. \end{aligned}$$

For $s^S < s^W$, we need

$$1 \geq \delta^{\tau-1}\Xi_1 + 1 - \delta^{\tau-1} + (1 - \delta)Q_1(1 - p)\frac{(1 - p)^{\tau-1} - \delta^{\tau-1}}{1 - p - \delta},$$

or, dividing by $\delta^{\tau-1}$,

$$1 - \Xi_1 \geq \frac{(1 - \delta)Q_1(1 - p)}{1 - p - \delta} \left[\frac{(1 - p)^{\tau-1}}{\delta^{\tau-1}} - 1 \right].$$

It is obvious that this will hold for small τ , where it reduces (for $\tau = 1$) to $1 - \Xi_1 \geq 0$. But $\frac{1-p}{\delta} > 1$, so that the left side grows without bound in τ , and so we will have $s^S > s^W$ for large τ .

We can take the limit as $\Delta \rightarrow 0$ and write this as

$$1 \geq Q_1 \frac{r}{r-p} (e^{(r-p)\tau} - 1).$$

This ensures that there is a lower interval of values of q for which $s^S(q) < s^W(q)$, but also a higher interval where this inequality is reversed. Hence, for $p < r$, each statement of the lemma holds.

Suppose instead that $p > 1 - \delta$ (or $p > r$). Then we have

$$\Xi_\tau = \delta^{\tau-1} \Xi_1 + 1 - \delta^{\tau-1} + (1-\delta)(1-p)Q_1 \delta^{\tau-2} \frac{1 - \left(\frac{1-p}{\delta}\right)^{\tau-1}}{1 - \frac{1-p}{\delta}} \quad (55)$$

$$= \delta^{\tau-1} \Xi_1 + 1 - \delta^{\tau-1} + (1-\delta)(1-p)Q_1 \delta^{\tau-1} \frac{1 - \left(\frac{1-p}{\delta}\right)^{\tau-1}}{\delta - (1-p)}. \quad (56)$$

For $s^S < s^W$, we need, for all τ ,

$$1 \geq \delta^{\tau-1} \Xi_1 + 1 - \delta^{\tau-1} + (1-\delta)(1-p)Q_1 \delta^{\tau-1} \frac{1 - \left(\frac{1-p}{\delta}\right)^{\tau-1}}{\delta - (1-p)}.$$

Dividing by $\delta^{\tau-1}$, this is

$$1 - \Xi_1 \geq (1-\delta)(1-p)Q_1 \frac{1 - \left(\frac{1-p}{\delta}\right)^{\tau-1}}{\delta - (1-p)}.$$

The problematic case is that in which τ gets arbitrarily large, giving

$$1 - \Xi_1 \geq (1-\delta)(1-p)Q_1 \frac{1}{\delta - (1-p)}.$$

This is

$$[\delta - (1-p)] \left(1 - \frac{1-\delta}{q_1}\right) \geq (1-\delta)(1-p) \frac{1-q_1}{q_1}.$$

To gain some insight here, multiply by q_1 and look at short time periods, making this

$$\Delta(p-r)(q_1 - r\Delta) \geq r\Delta(1-p\Delta)(1-q_1),$$

and hence, eliminating second-order terms

$$(p-r)q_1 \geq r(1-q_1),$$

or

$$pq_1 \geq r.$$

If this inequality holds, we will always have $s^S < s^W$. If it fails, we will again have a lower range of values of q with $s^S < s^W$, but an upper range where this inequality is reversed.

A.16 Proof of Lemma 11: Agent

Suppose we are in period τ , and hence there will be $\tau - 1$ additional belief revisions before rendering the players sufficiently pessimistic as to halt experimentation. Let z be the period, if any, in which only the more optimistic agent is induced to work.

Suppose first that the principal always induces both types of agent to work. Then using the definition of s^W , the payoff to the agent from working, given by

$$pq_\tau \pi(1 - s_\tau) + \delta(1 - pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}),$$

falls short of the payoff from shirking, given by

$$c + \delta \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}),$$

for all $s \in (s^W, s^S)$. On the other hand, suppose $z = \tau - 1$, so that in the next period using the definition of s^S , the agent's payoff from working, given by

$$pq_\tau \pi(1 - s) + \delta(1 - pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}),$$

exceeds that from shirking, given by

$$c + \delta W(\mathbf{1}_{q_\tau}, q_\tau),$$

for all $s \in (s^W, s^S)$. Fixing s , the payoff from working and the payoff from shirking are each upper-hemicontinuous, convex-valued (using the ability of the principal to mix) correspondences of z . There is accordingly for each $s \in (s^W, s^S)$ a value of z and a principal mixture that makes the agent indifferent.

A.17 Proof of Lemma 11: Principal

A.17.1 Outline

Suppose that the principal's belief assigns positive probability to two types. By Lemma 5, these are types that differ by one belief revision. In this subsection, we argue that it is optimal to have both types work for an initial number of periods, and then only one type work for one period, so that the posterior belief is degenerate afterwards. Of course,

the first phase might be either empty or take the entire horizon until the principal finds experimentation no longer profitable. And, crucially, there might be one period in which the principal is indifferent between having only the optimistic type or both work. The point is that the incentives to have both types work decrease over time (of course, once beliefs are degenerate, the agent works again.)

We refer to the event in which only the optimistic type works as a merger, since the public belief is degenerate afterwards. We consider three consecutive periods, and compare the relative value from merging in periods 0 and 1, with the relative value from merging in periods 1 and 2.

Our strategy is the following. We show that, if the principal prefers to merge in the third rather than in the second, she also prefers to merge in the second rather than the first, and hence she waits until some point before merging. This requires deriving first three payoffs for the principal, according to the period in which merger occurs. Then, we must compute the differences of consecutive values, and then compare those differences. In all three cases, since merger will have occurred within two periods, continuation payoffs to both players will be the same.

A.17.2 The Value of Merging

Before computing the relative values, we must compute the values from merging in each of the three periods. This requires solving a system of equations for the value of the agent (as a function of his belief and period), and the value to the principal, for each of the three cases.

Periods are labeled 0, 1, 2, which is identified by the first subscript in the notation. The second subscript refers to the agent's type: type k has belief q_k , where $q_0 > q_1 > q_2 > q_3$. So w_{tk} refers to the agent's payoff in period t with belief q_k . In period 0, the principal assigns probability μ to the agent having belief q_0 , and $1 - \mu$ to belief $q_1 = \varphi(q_0)$.

1. Let v^0 denote the value from merging in the first period (in period 0). In that case,

$$\begin{aligned} w_{00} &= (1 - s_0) pq_0 \pi + \delta (1 - pq_0) w_{11}, \\ w_{00} &= c + \delta \frac{q_0}{q_1} w_{11}, \\ w_{11} &= (1 - s_1) pq_1 \pi + \delta (1 - pq_1) w_{22}, \\ w_{11} &= c + \delta \frac{q_1}{q_2} w_{22}, \end{aligned}$$

as well as

$$\begin{aligned} w_{22} &= (1 - s_2) pq_2 \pi + \delta (1 - pq_2) w_3, \\ w_{22} &= c + \delta \frac{q_2}{q_3} w_3. \end{aligned}$$

We can solve this system, and plug the solution for shares in the principal's payoff. As a function of the continuation payoff v_3 (from the third period onward), this value is given by (before substituting)

$$v^0 = s_0 p q_0 \mu \pi - c + \delta [1 - p q_0 \mu] \times \\ (s_1 p q_1 \pi - c + \delta (1 - p q_1) (s_2 p q_2 \pi - c + \delta (1 - p q_2) v_3)).$$

2. Consider the system of equations in which merging (of beliefs) occurs in the second period. We have,

$$w_{12} = c + \delta w_{22},$$

as well as

$$w_{11} = c + \delta \frac{q_1}{q_2} w_{22}, \\ w_{11} = (1 - s_1) p q_1 \pi + \delta (1 - p q_1) w_{22},$$

and

$$w_{00} = (1 - s_0) p q_0 \pi + \delta (1 - p q_0) w_{11}, \\ w_{01} = (1 - s_0) p q_1 \pi + \delta (1 - p q_1) w_{12}, \\ w_{01} = c + \delta w_{11},$$

and finally

$$w_{22} = (1 - s_2) p q_2 \pi + \delta (1 - p q_2) w_3, \\ w_{22} = c + \delta \frac{q_2}{q_3} w_3.$$

Hence, the time-0 payoff to the principal from merging in the second period is

$$v^1 = s_0 p (\mu q_0 + (1 - \mu) q_1) \pi - c + \delta [1 - p (\mu q_0 + (1 - \mu) q_1)] \cdot \\ [s_1 p q_1 \mu_1 \pi - c + \delta (1 - p q_1 \mu_1) (s_2 p q_2 \pi - c + \delta (1 - p q_2) v_3)],$$

where

$$\mu_1 := \frac{\mu q_0}{\mu q_0 + (1 - \mu) q_1}.$$

Of course, beliefs satisfy, for $k = 0, 1, 2$:

$$q_{k+1} = \frac{(1 - p) q_k}{1 - p q_k}.$$

3. Finally, we consider the system in which two periods elapse before merging. This system is given by

$$\begin{aligned}
w_{00} &= (1 - s_0) p q_0 \pi + \delta (1 - p q_0) w_{11}, \\
w_{01} &= (1 - s_0) p q_1 \pi + \delta (1 - p q_1) w_{12}, \\
w_{11} &= (1 - s_1) p q_1 \pi + \delta (1 - p q_1) w_{22}, \\
w_{12} &= (1 - s_1) p q_2 \pi + \delta (1 - p q_2) w_{23}, \\
w_{23} &= c + \delta w_3, \\
w_{22} &= c + \delta \frac{q_2}{q_3} w_3, \\
w_{22} &= (1 - s_2) p q_2 \pi + \delta (1 - p q_2) w_3, \\
w_{12} &= c + \delta w_{22}, \\
w_{01} &= c + \delta w_{11},
\end{aligned}$$

and hence, solving this system, the payoff to the principal of this course of action is

$$\begin{aligned}
v^2 &= s_0 p (q_0 \mu + q_1 (1 - \mu)) \pi - c + \delta (1 - p (q_0 \mu + q_1 (1 - \mu))) \cdot \\
&\quad [s_1 p (q_1 \mu_1 + (1 - \mu_1) q_2) \pi - c + \delta (1 - p (q_1 \mu_1 + (1 - \mu_1) q_2)) v],
\end{aligned}$$

where

$$v := \mu_2 s_2 p q_2 \pi + \delta (1 - \mu_2 p q_2) v_3,$$

and

$$\mu_2 := \frac{\mu_1 q_1}{\mu_1 q_1 + (1 - \mu_1) q_2}.$$

Note that the (unknown) continuation payoffs v_3, w_3 after the third period are the same in all cases, as merging has occurred by then.

A.17.3 The Relative Value of Merging: Intuition

As mentioned, we are interested in the relationship between the differences

$$\Delta_1 := v^1 - v^0, \quad \Delta_2 := v^2 - v^1.$$

We need the result for a fixed (if arbitrarily small) $\Delta > 0$. However, because the argument is extremely tedious (see Section A.17.4), we first provide a simpler suggestive argument that holds in the limit, as $\Delta \rightarrow 0$.

We argue that $\Delta_1 < 0$ implies $\Delta_2 < 0$. This will be an easy consequence of the claim that $\Delta_1 = 0$ implies $\Delta_2 < 0$.

Let us take limits: $c \rightarrow c\Delta, p \rightarrow p\Delta, \delta \rightarrow 1 - r\Delta$, and of course $\Delta \rightarrow 0$. Fix μ . We get

$$\Delta_2 = (1 - \mu) [r(pq_1\pi - c) - \mu(1 - q_0)p[c + p(w_3 + q_0v_3 - q_0\pi)]] \Delta^2 + p\gamma_1\Delta^3 + o(\Delta^3),$$

where (letting $q := q_0, w := w_3$ and $v := v_3$ for notational simplicity),

$$\gamma_1 := (3(1 + q) - 2\mu(2 + q))rc - r(1 - \mu)(2q(2 + q)\pi - 3\mu(1 - q)(w + qv))p + \mu(1 - \mu)(1 - q)p[(1 + 5q)c - 3pq((1 + q)\pi - 2w - (1 + q)v)],$$

and

$$\Delta_2 = (1 - \mu) [r(pq\pi - c) - \mu(1 - q)p[c + p(w + qv - q\pi)]] \Delta^2 + p\gamma_2\Delta^3 + o(\Delta^3),$$

where

$$\gamma_2 := (4 + 5q - \mu(7 - \mu + q(2 + \mu)))rc - r(1 - \mu)(q\pi(7 - \mu + q(2 + \mu)) - 2\mu(1 - q)(w + qv))p + \mu(1 - \mu)(1 - q)p[(3 + 4q)c + (2w - q((5 + 2q)\pi - 2qv - 5(w + v)))] .$$

These expressions might be positive or negative, depending (among other things) on v, w . Note that there is no first-order term, which should not be surprising, given that we are considering differences. Note also that the second-order term is identical, which means we must look at the third-order term, namely γ_2 vs. γ_1 .

We will solve for the value of μ for which $\Delta_1 = 0$, and plug it in our expression for Δ_2 , showing that $\Delta_2 < 0$, and hence $\Delta_2 < \Delta_1$. The general result will then follow from monotonicity of these expressions in μ .

Note that $\gamma_1 > 0$ for $\mu = 0$, and $\gamma_1 < 0$ for $\mu = 1$. It is not hard to show that Δ_1 is monotone in μ , so that there exists at most one solution $\mu \in [0, 1]$ to $\Delta_1 = 0$. It is also easy to see that $\mu \rightarrow 1$ as $\Delta \rightarrow 0$. We claim that this solution must be given by

$$\mu = 1 - \kappa\Delta + o(\Delta),$$

for some κ to be defined. Indeed, plugging in such an expression into Δ_1 and solving for the root gives us

$$\kappa = \frac{(1 - q)rcp}{r(pq\pi - c) - p(1 - q)(c + p(w + qv - q\pi))}.$$

For our claim to make sense, we must check that $\kappa > 0$ (so that our formula gives us a well-defined probability).

Claim A.1 *It holds that $\kappa > 0$.*

Proof of claim: We have to consider different cases. Throughout, we conserve on clutter by normalizing c to 1.

First, consider the high surplus, impatient case, in which $a > 4$ and $a - \sigma > 2$ (where $a := p\pi > 2$). We insert the closed-form formulas for v, w to get

$$\kappa = \frac{1 - \sigma^2}{(1 - q) q^2 \sigma^2 (a^2 + \sigma (8 - a (8 - a))) \left(\frac{2(1-q)}{q(a-2)} \right)^{1+\frac{1}{\sigma}} - 2B},$$

where

$$B := 1 + \sigma - q \left[\sigma (1 + 3\sigma - 2q\sigma + 2(1 - q)^2 \sigma) + a(1 - \sigma) (1 + \sigma (2 - q + \sigma (1 - q)^2)) \right].$$

We use the Bernoulli inequality to bound the term $\left(\frac{2(1-q)}{q(a-2)} \right)^{\frac{1}{\sigma}}$ that appears in the denominator, noting that the (upper or lower) bound it provides (according to whether $\sigma \leq 1$) gives us a lower bound to μ , given that the numerator changes signs at $\sigma = 1$ as well. The resulting expression is positive if and only if

$$\begin{aligned} S_0(\sigma) &:= a^3 q (2 + \sigma (3 + q^2 (\sigma - 1) + \sigma (1 - 2q))) \\ &\quad - 2a^2 (1 + \sigma + 4q + \sigma q (9 - 5q + (1 - q)^2)) \\ &\quad - 8 - 8\sigma (2 - q) + 8a (1 + q) (1 + \sigma (2 - q)) > 0. \end{aligned}$$

We claim that S_0 is increasing. First, it is convex, as its second derivative is

$$2(a - 2) a^2 (1 - q)^2 q > 0.$$

Second, its first derivative, evaluated at 0, equals

$$\frac{4(a - 1)(a - 2)^2}{a} + 3(a - 2)^3 \tilde{q} + 4(a - 2) a \tilde{q}^2 - a^3 \tilde{q}^3,$$

where $\tilde{q} = q - \underline{q}$. This is positive for all $\tilde{q} < 1 - a$ if it is positive for $\tilde{q} = 1 - a$, and for $\tilde{q} = 1 - a$, we obtain an expression which is increasing in a and equal to 0, for $a = 2$ —so, positive everywhere.

We are left to claim that $S_0(0) > 0$. But

$$S_0(0) = 2(2 - a)^2 (aq - 1) > 0,$$

since $q > \underline{q} = 2/a$. This establishes that $\kappa > 0$ in case 1.

Next, suppose we have a high surplus, patient project. Note that nothing in our previous analysis hinged on patience, so the claim is true for $q < q^*$. So we can focus on the case in which the principal's profit is zero. In that case, we know that

$$v = 0, w = q\pi - \frac{2c}{p},$$

and inserting into κ , we get

$$\kappa = \frac{(1-q)r\sigma}{aq-1+(1-q)\sigma} > 0.$$

This also applies to the other cases whenever there is zero profit to the principal. This establishes that $\kappa > 0$ in case 2.

We are left with the case of low surplus and impatience, when the belief is such that the principal's payoff is positive, i.e. $q > q^{**}$. Using the boundary conditions $w(q^{**}) = q^{**}\pi - \frac{2c}{p}$, $v(q^{**}) = 0$, and our formula for q^{**} , we can solve for the two differential equations that give v, w and plugging into κ gives that κ is of the sign of

$$aq-1 + \frac{\sigma(1-q)}{1-\sigma^2} [q(1-(1-q)\sigma)(a(1-\sigma)+2\sigma) - 1 - \sigma + C],$$

where

$$C := (1-q)q\sigma(a+(a-6)\sigma) \left(\frac{1-q}{q} \frac{\sigma-2}{a-2-\sigma} \right)^{\frac{1}{\sigma}}.$$

Again, we note that using Bernoulli's inequality, applied to C , provides us a lower bound to this expression whether or not σ is less than one. Simplifying gives us the following expression:

$$S_1(q) := aq-1 + \frac{(1-q)\sigma(2+a^2q^2+\sigma(8q-7-4q(1-q)\sigma)+a(1+\sigma+2q(\sigma-2-2q\sigma)))}{(a-2-\sigma)(1+\sigma)}.$$

First, we have

$$S_1'''(q) = -\frac{6(a-2\sigma)^2\sigma}{(a-2-\sigma)(1+\sigma)} < 0.$$

Next,

$$S_1''(q^{**}) = \frac{2(a-2\sigma)\sigma}{\sigma+1} > 0,$$

while

$$S_1''(1) = -\frac{4(a-2\sigma)\sigma}{\sigma+1} < 0.$$

Hence, S'' is first convex, then concave. Note that

$$S_1(q^{**}) = \frac{a - \sigma^2}{a - 2\sigma} > 0,$$

and

$$S_1(1) = a - 1 > 0,$$

and finally that

$$S_1'(q^{**}) = a - \sigma > 0.$$

Hence, $S_1(q) > 0$ for all $q \in [q^{**}, 1]$, and we are done with this case as well.

We have now verified that $\kappa > 0$ in all cases, and that our expansion for the root μ of Δ_1 is valid. **End of the proof of claim.**

We now can come back to our comparison between Δ_1 and Δ_2 . Plugging in our formula for κ (and hence μ) into Δ_2 , we get

$$\Delta_2 = -(1 - q)rcp\Delta^3 + o(\Delta^3) < 0.$$

Because Δ_2 is also monotone in μ , it follows that, more generally,

$$\Delta_1 < 0 \implies \Delta_2 < 0,$$

which establishes “concavity:” if the principal is indifferent between merging and not in some period, she strictly prefers to merge after, and strictly prefers to keep beliefs separate before, for all $\Delta > 0$ small enough.

A.17.4 The Relative Value of Merging: Formal Analysis

The remainder of this section provides a formal analysis for fixed (non-vanishing) $\Delta > 0$. The strategy of proof is exactly the same: setting $\Delta_1 = 0$, solving for one of the “parameters” (in this case, v), and showing that $\Delta_2 < 0$ for that value. Unfortunately, as mentioned, the analysis is significantly more tedious. In what follows, we drop the reference to Δ , keeping in mind that it is fixed once and for all, to a value that is such that $\delta = \exp(-r\Delta) \geq 2/3$, and $p\Delta \leq 1/2$.

Now, using the formulas for payoffs, we have $\Delta_1 \propto A_1 - A_2$, where

$$A_1 := p(\mu - 1) \left[\pi(p - 1)^3 q (p(q - 1)\mu (pq((p - 1)^2 \delta^2 - p + 2) - 1) + (p - 1)(\delta - 1)((p - 2)pq + 1)) \right. \\ \left. - p(q - 1)\delta^3 \mu (p((p - 3)p + 3)q - 1) ((p - 1)^3 qv + p((p - 3)p + 3)qw - w) \right],$$

and

$$A_2 := c(p - 1)((p - 2)pq + 1) [p^2(q - 1)\delta^2(\mu - 1)\mu((p - 2)p + 2)q - 1 - \delta(pq - 1)(p^2(q - 1)^2\mu^2 + (p - 1)\mu(p(q - 2) + 1) + (p - 1)^2) \\ + (pq - 1)(p(q - 2)\mu + p + \mu - 1)(p(q - 1)\mu + p - 1)]$$

It holds that Δ_1 is decreasing in v : its derivative w.r.t. v is

$$-\delta^3(p-1)^3p^2(q-1)q(\mu-1)\mu(p((p-3)p+3)q-1) \leq 0.$$

Similarly, Δ_2 is decreasing in v , as the second derivative is

$$\delta^2(p-1)^5p^2(q-1)q(\mu-1)\mu(pq(p((q-1)\mu+1)-2)+1) \leq 0.$$

Hence, it suffices to solve for v such that $\Delta_1 = 0$ and show that $\Delta_2 \leq 0$ for that value of v . Solving gives $v = (A_4 + A_5)/A_3$, where

$$A_3 := -(p-1)^3p^2(q-1)q\delta^3(\mu-1)\mu(p((p-3)p+3)q-1),$$

and

$$A_4 := -p(\mu-1) \left[\pi(p-1)^3q(p(q-1)\mu(pq((p-1)^2\delta^2-p+2)-1) + (p-1)(\delta-1)((p-2)pq+1)) \right. \\ \left. -p(q-1)w\delta^3\mu(p((p-3)p+3)q-1)^2 \right],$$

and

$$A_5 := c(p-1)((p-2)pq+1) \left[p^2(q-1)\delta^2(\mu-1)\mu((p-2)p+2)q-1 - \delta(pq-1)(p^2(q-1)^2\mu^2 + (p-1)\mu(p(q-2)+1) + (p-1)^2) \right. \\ \left. + (pq-1)(p(q-2)\mu+p+\mu-1)(p(q-1)\mu+p-1) \right].$$

We now insert the value of v in Δ_2 and get that $\Delta_2 = \Delta^* := \frac{\delta-1}{\delta(p((p-3)p+3)q-1)}(B_1 + c((p-2)pq+1)(B_2 + \delta B_3))$, with

$$B_1 := \pi(p-1)^5pq(1-\mu)(pq(p(q-1)\mu+p-2)+1) \left[(p-1)\delta(p((q-1)\mu((p-2)^2pq-1) + (p-1)((p-3)p+3)q-1) + 1) \right. \\ \left. + ((p-2)pq+1)(p(-q\mu+\mu-1)+1) \right],$$

and

$$B_2 := (p-1)^3(pq-1)(p(q-2)\mu+p+\mu-1)(p(q-1)\mu+p-1)(pq(p(q-1)\mu+p-2)+1),$$

and finally $B_3 = C_1 + C_2 + C_3 + C_4$, with

$$C_1 := p^3(q-1)^2\mu^3((p-2)^2pq-1) \left(p(((p-3)p+3)q^2 + (p((p-5)p+10) - 9)q - p + 2) + q \right),$$

as well as

$$C_2 := (1-p)p(q-1)\mu^2 \left[p \left[p \left[p(p(p((p-9)p+34) - 69) + 75) - 35 \right] q^3 \right. \right. \\ \left. \left. - (p-2)(p(p(p(3p-19) + 53) - 74) + 44) + 2 \right) q^2 + (p-2)(p(p+2)(2p-5) + 24)q - 3p + 3 \right] + 6q + 4 \right] - 1,$$

and

$$C_3 := (p-1)^2\mu \left[p \left[p \left[p(p(p((15-2p)p-49) + 85) - 79) + 31 \right] q^3 \right. \right. \\ \left. \left. + (p(p(p(p(3p-23) + 78) - 139) + 128) - 39) - 11 \right] q^2 - p(p(p(p+2) - 21) + 50)q + 2p + 38q + 1 \right] - 3q - 5 \right] + 1,$$

and finally

$$C_4 := -(p-1)^4(p((p-3)p+3)q-1)^2.$$

We normalize throughout c to 1. We now note that

$$\begin{aligned} \frac{d\Delta^*}{d\pi} &= (p-1)^5 pq(1-\mu)(pq(p(q-1)\mu+p-2)+1) [(p-1)\delta(p((q-1)\mu((p-2)^2 pq-1)+(p-1)((p-3)p+3)q-1)+1) \\ &\quad +((p-2)pq+1)(p(-q\mu+\mu-1)+1)], \end{aligned}$$

which is always negative. To see this, note that $(p-1)^5 pq(1-\mu)(pq(p(q-1)\mu+p-2)+1) < 0$. The last factor appearing in $\frac{d\Delta^*}{d\pi}$ is linear in μ , with coefficient equal to $p(q-1)((p-2)pq(d(p-2)(p-1)-1)+d(-p)+d-1) > 0$, i.e., it is increasing in μ ; yet at $\mu=0$ it equals $(p-1)(\delta(p-1)(p((p-3)p+3)q-1)-(p-2)pq-1) > 0$, and so it is positive for all $\mu \in [0, 1]$.

To check that $\Delta^* \leq 0$, it thus suffices to consider the case in which $\pi = 2/(pq)$. We also note that $\Delta^* = \Delta^*(\delta)$ is linear in δ , and so it suffices to consider the two extreme cases $\delta = 1$ and $\delta = 2/3$. In the former case, we get $\Delta^*(1) = D_1 + p((p-2)pq+1)(D_2+D_3+D_4+D_5)$, where

$$D_1 := 2(p-1)^5 p(1-\mu)(pq(p(q-1)\mu+p-2)+1)(p((q-1)\mu((p-2)((p-3)p+1)q-1)+(p-2)((p-3)p+3)q-1)+q+1),$$

$$\begin{aligned} D_2 &:= p^3(q-1)^2 \mu^3 [p((p-3)p((p-3)p+7)+11)q^3 \\ &\quad + (p(p(p((p-9)p+32)-63)+70)-35)+1)q^2 - (p-4)((p-2)(p-1)p+1)q+p-2], \end{aligned}$$

$$\begin{aligned} D_3 &:= (1-p)p(q-1)\mu^2 [p [p(p(p((p-9)p+31)-59)+64)-31)q^3 \\ &\quad + (p(p(p((25-3p)p-86)+164)-177)+84)+2)q^2 + (p(p(2(p-5)p+13)+14)-31)q-5p+8] + 3q], \end{aligned}$$

as well as

$$D_4 := (p-1)^4(q(p(-((p-3)p((p-3)p+5)+7)q+p-2)+3)-1),$$

and finally

$$\begin{aligned} D_5 &:= -(p-1)^2 \mu [p [p(p(p(p(2p-15)+46)-74)+66)-26)q^3 \\ &\quad + (p(p(p((23-3p)p-74)+126)-119)+44)+6)q^2 + (p(p((p-2)p-2)+23)-25)q-5p+6] + 2q]. \end{aligned}$$

It is a matter of tedious algebra to show that this polynomial function is always negative for $p \leq 1/2$. Here are the details. First, we show that $\Delta^*(1)$ is concave in q . As $d^2\Delta^*(1)/dq^2$ is equal to

$$\begin{aligned} &-2(1-p)^6 p^2 [\mu^3(p-3)(p-1)p^2 + \mu^2((p-4)(p-3)(p-1)p+1) \\ &\quad + \mu(p(p((13-3p)p-18)+4)+2) + p(p((p-6)p+14)-16)+9] \leq 0 \end{aligned}$$

(leaving aside the negative factor $-2(1-p)^6 p^2$, the remainder is decreasing in p , yet positive at $p = 1/2$), at $q = 1$, it is enough to show that its third derivative, $\Delta^{(3)}(1)$

is positive (so that $d^2\Delta^*(1)/dq^2$ is increasing.) In turn, because $\Delta^{(3)}(1)$, evaluated at 1, equals

$$-6(1-p)^4p^3 [(p-2)(p-1)p((p-4)p+2)\mu^3 - (p-1)(p((p-2)p(3p-5)+6)-6)\mu^2 + p(p(p((p-9)p+31)-58)+62)\mu + p(p(p((p-8)p+26)-43)+37)-14(2\mu+1)] \geq 0,$$

(leaving aside the negative factor $-6(1-p)^4p^3$, this is decreasing in μ and increasing in p , yet negative at $(\mu, p) = (0, 1/2)$), it suffices to argue that the fourth derivative w.r.t. q , $\Delta^{(4)}(1)$, is negative. To show this, we argue that $\Delta^{(5)}(1)$ is positive, yet $\Delta^{(4)}(1)$, evaluated at 1, is

$$24(1-p)^2p^4\mu [(p-1)^2(p^4-3p^3+6p-3)\mu^2 + p(p(p(p^3-4p^2+p+30)-82)+85)\mu + p(p(p(p(19-2p)p-76)+166)-214)+158]-31\mu-52] \leq 0.$$

(Leaving aside the positive factor $24(1-p)^2p^4\mu$, this is decreasing in μ and increasing in p , yet negative at $(p, \mu) = (1/2, 0)$). In turn, to show that $\Delta^{(5)}(1) \geq 0$, we note that the sixth derivative is

$$720(p-2)p^6((p-3)p((p-3)p+7)+11)\mu^3 \leq 0,$$

and that, evaluated at 1, the fifth derivative is

$$120(p-1)p^5\mu^2(p(p(p(p(p(\mu-1)-6\mu+11)+12\mu-49)-2\mu+121)-2(12\mu+91))+29\mu+159)-9\mu-62) \geq 0.$$

(Leaving aside the negative factor $120(p-1)p^5\mu^2$, this is increasing in p and negative at $p = 1/2$). Having shown that $\Delta^*(1)$ is concave in q , we note that $d\Delta^*(1)/dq$ evaluated at $q = 1$, is

$$(1-p)^8p(1-p(\mu^2(2(p-3)p+3)+\mu(-p^2+p+3)-(p-2)^2)) \geq 0,$$

(leaving aside the positive factor $(1-p)^8p$, the last factor is decreasing in μ and equals $(1-p)^2$ at $\mu = 1$), so it is increasing in q . Yet it is equal to $-(1-\mu)(2-p)p(1-p)^{10} < 0$ at $q = 1$, so it is negative everywhere.

Similarly, for $\delta = 2/3$, we get $\Delta^*(2/3) = E_1 + ((p-2)pq+1)(E_2 + 2/3(E_3 + E_4 + E_5))$, where

$$E_1 := -\frac{2}{3}(\mu-1)(p-1)^5(pq(p(\mu(q-1)+1)-2)+1)(p(\mu(q-1)((p-2)(2(p-3)p+1)pq-2p-1)+(p-1)(2(p-4)p+9)pq-2p+1)+1),$$

$$E_2 := (p-1)^3(pq-1)(\mu(p(q-2)+1)+p-1)(p(\mu(q-1)+1)-1)(pq(p(\mu(q-1)+1)-2)+1),$$

$$E_3 := \mu^3p^3(q-1)^2((p-2)^2pq-1)(p(((p-3)p+3)q^2+(p((p-5)p+10)-9)q-p+2)+q)-(p-1)^4(p((p-3)p+3)q-1)^2,$$

$$E_4 := (1-p)p(q-1)\mu^2[p[p[p(p(p((p-9)p+34)-69)+75)-35]q^3 - (p-2)(p(p(p(3p-19)+53)-74)+44)+2]q^2 + (p-2)(p(p+2)(2p-5)+24)q - 3p+3) + 6q+4) - 1],$$

and

$$E_5 := (p-1)^2 \mu [p [p [p(p(p((15-2p)p-49)+85)-79)+31)q^3 \\ + (p(p(p(p(3p-23)+78)-139)+128)-39)-11)q^2 - p(p(p(p+2)-21)+50)q + 2p + 38q + 1) - 3q - 5) + 1].$$

Again, it is tedious but straightforward to show that this is negative provided $p \leq 1/2$ (the steps are the same, consisting in showing that the derivatives w.r.t. q alternate in sign). The case in which there is delay is dealt with similarly (solving for w from $\Delta_1 = 0$, inserting into Δ_2 and showing it is negative).

A.18 Proof of Lemma 12

Refer to each passage of Δ time as a period, and let periods be numbered beginning with 0, with subsequent periods numbered 1, 2, ... Suppose the principal makes an offer $s \in (s^W, s^S)$ in period 0, facing agent belief q_0 . The agent mixes between working and shirking. The principal then makes offers that induce both the optimistic and pessimistic agent to work, until reaching period t , in which only the optimistic agent works. The agents' beliefs are then "merged," after which both agents work in each period. We argue that this gives the principal a lower payoff than does s^W .

If the principal makes the equilibrium offer s^W in period 0, she receives

$$V(\mathbf{1}_{q_0}, q_0) = pq_0\pi s^W - c + \delta(1 - pq_0)V(\mathbf{1}_{q_1}, q_1).$$

If the principal instead makes offer $s \in (s^W, s^S)$ and the agent shirks, then the principal's payoff is

$$-c + \delta\tilde{V},$$

for some continuation payoff \tilde{V} that satisfies

$$\tilde{V} \leq V(\mathbf{1}_{q_0}, q_0).$$

The latter inequality follows from the observations that the (i) continuation payoffs $V(\mathbf{1}_{q_0}, q_0)$ and \tilde{V} are both generated by continuation paths under which the principal makes some number T of offers, to posteriors q_0, q_1, \dots, q_{T-1} , (ii) along both continuation paths, the agent works in every period, (iii) the principal's offers under continuation paths $V(\mathbf{1}_{q_0}, q_0)$ and \tilde{V} from the t -th period on are identical under the two paths, and (iv) during the first $t-1$ period, the outcome generating payoff \tilde{V} faces an additional set of constraints not imposed on $V(\mathbf{1}_{q_0}, q_0)$, namely that the more pessimistic agent also be willing to work.

Hence, the principal's payoff is lower under s if the agent shirks. Suppose instead the principal offers s and the agent works. We again need to show that this generates a lower payoff for the principal than offering s^W . The continuation paths following s^W and s generate identical outcomes over the periods 1, ..., $t-1$. The path following s^W

then continues with payoff $V(\mathbf{1}_{q_t}, q_t)$, while the path following s delays this payoff by one period of shirking. Hence, we need to show that

$$pq_0\pi s + \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)(-c + \delta V(\mathbf{1}_{q_t}, q_t)) < pq_0\pi s^W + \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)V(\mathbf{1}_{q_t}, q_t),$$

or

$$pq_0\pi(s - s^W) < \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)[c + (1 - \delta)V(\mathbf{1}_{q_t}, q_t)].$$

We consider the worst case in terms of satisfying this inequality, namely that in which $V(\mathbf{1}_{q_t}, q_t) = 0$, and hence we need

$$pq_0\pi(s - s^W) < \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)c. \quad (57)$$

In response to offer s^W , the agent is indifferent between shirking and working. We let W_t denote the continuation payoff received by the agent after period t , along the Markov-equilibrium path. We let $W_{1,t-1}$ identify the payoffs collected by the agent between periods 1 and $t-1$ (inclusive), along this equilibrium path. Then the condition for the agent to be indifferent is

$$pq_0\pi(1-s^W) + \delta(1-pq_0) \left[W_{1,t-1} + \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t \right] = c + \delta \frac{q_0}{q_1} \left[W_{1,t-1} + \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t \right]. \quad (58)$$

Under offer s , a working agent receives payoff

$$pq_0\pi(1-s) + \delta(1-pq_0) \left[W_{1,t-1} + \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)(c + \delta W_t) \right],$$

while a shirking agent receives

$$c + \delta \frac{q_0}{q_1} W_{1,t-1} + \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1}. \quad (59)$$

We can use (58) to rewrite the payoff (59) of a shirking agent as

$$\begin{aligned} & pq_0\pi(1-s^W) + \delta(1-pq_0) \left[W_{1,t-1} + \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t \right] \\ & - \delta \frac{q_0}{q_1} \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t + \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1}. \end{aligned}$$

The condition that the agent be indifferent between shirking and working, after offer s , is then

$$\begin{aligned} & -pq_0\pi s + \delta(1-pq_0) \left[\delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)c + \delta^t \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t \right] \\ = & -pq_0\pi s^W + \delta(1-pq_0)\delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t - \delta \frac{q_0}{q_1} \delta^{t-1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t + \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1}. \end{aligned}$$

We can rewrite this as

$$\begin{aligned} pq_0\pi(s - s^W) &= \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)c + \delta^{t+1} \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t \\ &\quad + \delta^t \frac{q_0}{q_1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1}. \end{aligned}$$

Hence, to establish the result, we need to show that

$$\delta^{t+1} \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t + \delta^t \frac{q_0}{q_1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1} < 0.$$

Eliminating some common terms, this is

$$\delta W_t - W_t + \frac{q_0}{q_1} \frac{1}{1 - pq_0} W_t - \frac{W_{t-1}}{1 - pq_{t-1}} < 0.$$

Successive manipulations now give

$$\begin{aligned} \delta W_t - W_t + \frac{W_t}{1 - p} - \frac{W_{t-1}}{1 - pq_{t-1}} &< 0, \\ \frac{1 - pq_{t-1}}{1 - p} W_t - \left(c + \delta \frac{q_{t-1}}{q_t} W_t \right) &< (1 - \delta)(1 - pq_{t-1})W_t, \\ \frac{q_{t-1}}{q_t} W_t - \delta \frac{q_{t-1}}{q_t} W_t &< c + (1 - \delta)(1 - pq_{t-1})W_t, \\ (1 - \delta) \frac{q_{t-1}}{q_t} W_t &< c + (1 - \delta)(1 - pq_{t-1})W_t, \\ (1 - \delta) \left(\frac{q_{t-1}}{q_t} - (1 - pq_{t-1}) \right) W_t &< c, \\ (1 - \delta) \frac{q_{t-1}}{q_t} \left(1 - \frac{q_t}{q_{t-1}} (1 - pq_{t-1}) \right) W_t &< c, \\ (1 - \delta) \frac{q_{t-1}}{q_t} p W_t &< c, \end{aligned}$$

which holds for small Δ .

A.19 Proof of Lemma 13

We have four cases to consider.

First, if $\tilde{s}^S \leq \tilde{s}^W$ and $s^S \leq s^W$, the result is immediate. The agents do not mix in these circumstances, and the optimality of the principal's strategy follows from the fact that the agent whose belief is (\tilde{q}) (alternatively, q) shirks if and only if the offer falls short of \tilde{s}^W (or s^W).

Second, suppose $\tilde{s}^S \geq \tilde{s}^W$ and $s^S \leq s^W$. Then we need to consider offers in $(\tilde{s}^W, \tilde{s}^S)$. The argument follows that of the fourth case below.

Third, suppose $\tilde{s}^S \leq \tilde{s}^W$ and $s^S \geq s^W$. We must then consider offers in (s^W, s^S) . Here, the result is straightforward. For any such offer, the agent believing \tilde{q} shirks. The payoff and the continuation play, if the agent believes \tilde{q} , is then independent of the current offer, and we can condition on the event that the agent believes q . Here, Lemma 12 uses only the information that this agent is indifferent between working and shirking to show that, no matter what the agent's action, the principal receives a lower payoff than would be the case under offer s^W .

We thus have the case $\tilde{s}^S \geq \tilde{s}^W$ and $s^S \geq s^W$. An argument identical to that of the preceding case addresses offers $s \in (s^W, s^S)$. This allows us to focus attention on the case $s \in (\tilde{s}^W, \tilde{s}^S)$. Here, the \tilde{q} agent mixes between working and shirking, while the q agent works. If the agent happens to have belief \tilde{q} , then the proof of Lemma 12 applies to ensure that the principal is better off with offer \tilde{s}^W than s , whether the agent shirks or works. If the agent is type q , then we need to establish the counterpart of (57), or

$$pq_{-1}\pi(s - \tilde{s}^W) < \delta^t \prod_{\tau=-1}^{t-2} (1 - pq_\tau)c,$$

where we take $q = q_{-1}$ and $\tilde{q} = q_0$. The information we have available is that the agent believing $\tilde{q} = q_0$ is indifferent, or

$$\begin{aligned} pq_0\pi(s - s^W) &= \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)c + \delta^{t+1} \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t \\ &\quad + \delta^t \frac{q_0}{q_1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1}. \end{aligned}$$

Hence, to establish the result, we need to show that

$$\begin{aligned} &\delta^{t+1} \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)W_t + \delta^t \frac{q_0}{q_1} \prod_{\tau=1}^{t-1} (1 - pq_\tau)W_t - \delta^t \prod_{\tau=0}^{t-2} (1 - pq_\tau)W_{t-1} \\ &< \frac{q_0}{q_{-1}} \delta^t \prod_{\tau=-1}^{t-2} (1 - pq_\tau)c - \delta^t \prod_{\tau=0}^{t-1} (1 - pq_\tau)c. \end{aligned}$$

Eliminating some common terms, this is

$$\delta W_t - W_t + \frac{q_0}{q_1} \frac{1}{1 - pq_0} W_t - \frac{W_{t-1}}{1 - pq_{t-1}} < c \left(\frac{q_0}{q_{-1}} \frac{1 - pq_{-1}}{1 - pq_{t-1}} - 1 \right).$$

Successive manipulations now give

$$\begin{aligned}
\delta W_t - W_t + \frac{W_t}{1-p} - \frac{W_{t-1}}{1-pq_{t-1}} &< c \left(\frac{1-p}{1-pq_{t-1}} - 1 \right), \\
\frac{1-pq_{t-1}}{1-p} W_t - \left(c + \delta \frac{q_{t-1}}{q_t} W_t \right) - (1-\delta)(1-pq_{t-1})W_t &< cp(q_{t-1}-1), \\
\frac{q_{t-1}}{q_t} W_t - \delta \frac{q_{t-1}}{q_t} W_t - c - (1-\delta)(1-pq_{t-1})W_t &< cp(q_{t-1}-1), \\
(1-\delta) \frac{q_{t-1}}{q_t} W_t - c - (1-\delta)(1-pq_{t-1})W_t &< cp(q_{t-1}-1), \\
(1-\delta) \left(\frac{q_{t-1}}{q_t} - (1-pq_{t-1}) \right) W_t - c &< cp(q_{t-1}-1), \\
(1-\delta) \frac{q_{t-1}}{q_t} \left(1 - \frac{q_t}{q_{t-1}}(1-pq_{t-1}) \right) W_t - c &< cp(q_{t-1}-1), \\
(1-\delta) \frac{q_{t-1}}{q_t} p W_t - c &< cp(q_{t-1}-1),
\end{aligned}$$

which holds for small Δ .

A.20 Proof of Lemma 14

We now examine the smallest payoff available to the principal in the final period of a no-delay equilibrium. The task is to minimize s_1 . To do this, we assume that should the principal choose a larger value of s_1 , the agent is expected to shirk. We are thus identifying the value s^S , via the following constraint:

$$c + \delta pq_1 \pi(1-s_0) \geq pq_1 \pi(1-s_0) + \delta(1-pq_1) \max\{c, pq_0 \pi(1-s_0)\}.$$

The Markov restriction will require setting $s_0 = s_1$. Notice, however, that if we are allowed to set these separately, then minimizing s_1 is achieved by minimizing s_0 . Hence, the minimum final-period principal payoff, over all no-delay equilibria, can be achieved by a Markov equilibrium. Hence, we can write

$$c + \delta pq_1 \pi(1-s) \geq pq_1 \pi(1-s) + \delta(1-pq_1) \max\{c, pq_0 \pi(1-s)\}. \quad (60)$$

We now argue that for q_1 sufficiently close to \underline{q} , we can set the principal's payoff equal to zero. This requires showing that s with $pq_1 \pi s = c$ can satisfy the incentive constraint (60). First, we notice that for q_1 close to \underline{q} , we have $c > pq_0 \pi(1-s)$.³² Hence, using $pq_1 \pi s = c$, from (60) we need to show

$$(2-\delta)c \geq (1-\delta)pq_1 \pi + \delta(1-pq_1)c,$$

³²We need to show $c > pq_0 \pi(1-s)$, or $c > \frac{q_0}{q_1} pq_1 \pi(1-s) = \frac{q_0}{q_1} (pq_1 \pi - c)$ (using $pq_1 \pi s = c$). For q_1 very close to \underline{q} , this is $c \geq \frac{q_0}{q_1} (2c - c)$, which holds.

or

$$2(1 - \delta)c \geq (1 - \delta)(pq_1\pi - c).$$

But for q_1 sufficiently close to \underline{q} , we have $pq_1\pi - c$ arbitrarily close to c , giving the result. Hence, for $q_1 \in [\underline{q}, \hat{q}]$ for some \hat{q} , the principal's lowest Markov equilibrium payoff is 0.

Now let us examine values of q_1 large enough that q_0 is very close to \underline{q} . Here, we have $c < pq_0\pi(1 - s)$.³³ We now show that we cannot reduce the principal's payoff to zero in this case. This is equivalent to showing that we cannot satisfy (60) with $pq_1\pi s = c$, or equivalently that it is impossible that (with subsequent simplifications)

$$\begin{aligned} c + \delta pq_1\pi(1 - s) &\geq pq_1\pi(1 - s) + \delta(1 - pq_1)pq_0\pi(1 - s), \\ (2 - \delta)c &\geq (1 - \delta)pq_1\pi + \delta(1 - pq_1)pq_0\pi(1 - s), \\ (2 - \delta)c &\geq (1 - \delta)\frac{q_1}{q_0}2c + \delta(1 - pq_1)\left(2c - \frac{q_0}{q_1}c\right), \\ -\delta &\geq -\delta\frac{q_1}{q_0}2 + \delta(1 - pq_1)\left(2 - \frac{q_0}{q_1}\right), \\ -\delta &\geq -\delta\frac{q_1}{q_0}2 + \delta\left(2(1 - p)\frac{q_1}{q_0} - q_1\right), \\ -\delta q_0 &\geq -2\delta q_1 + 2\delta(1 - p)q_1 - \delta q_0 q_1, \\ -\delta q_0 &\geq -2\delta pq_1 - \delta q_0 q_1, \\ \delta q_0 &\leq 2\delta pq_1 + \delta q_0 q_1, \end{aligned}$$

which fails for small Δ (and hence small p).

To calculate the principal's minimum payoff over the region $[\hat{q}, \tilde{q}]$, we note that the principal's payoff is given by $pq_1\pi s - c$ for the lowest value of s , which satisfies the agent's incentive constraint, given by

$$c + \delta pq_1\pi(1 - s) = pq_1\pi(1 - s) + \delta(1 - pq_1)pq_0\pi(1 - s).$$

³³Using the fact that q_0 is set at its upper limit if \underline{q} and hence $pq_0\pi = 2c$, we have

$$\begin{aligned} c &\leq pq_0\pi(1 - s) \\ &= \frac{q_0}{q_1}pq_1\pi(1 - s) \\ &= pq_0\pi - \frac{q_0}{q_1}c \\ &= 2c - \frac{q_0}{q_1}c. \end{aligned}$$

Successive manipulations give

$$\begin{aligned}
c &= (1 - \delta)pq_1\pi(1 - s) + \delta(1 - pq_1)pq_0\pi(1 - s), \\
c &= (1 - \delta)pq_1\pi(1 - s) + \delta(1 - pq_1)pq_0\pi(1 - s), \\
c &= (1 - \delta)pq_1\pi(1 - s) + \delta(1 - pq_1)\frac{q_0}{q_1}pq_1\pi(1 - s), \\
c &= pq_1\pi(1 - s)[(1 - \delta) + \delta(1 - p)],
\end{aligned}$$

and hence

$$pq_1\pi s = pq_1\pi - \frac{c}{1 - \delta p},$$

with a principal's payoff of

$$pq_1\pi - \frac{2 - \delta p}{1 - \delta p}c.$$

A.21 Proof of Lemma 15

Fix a posterior q . Let $W^*(\mathbf{1}_q, q)$ be the agent's value in the no-delay principal-optimum Markov equilibrium, given that the agent holds belief q and the principal holds a degenerate belief concentrated on the value q . We seek a lower bound on the principal's payoff in any equilibrium. The strategy of proof is to note that one feasible option for the principal is to induce the agent to work in every period. Then we ask what is the most expensive such a strategy could be for the principal, or equivalently, what is the largest equilibrium payoff for the agent in an equilibrium in which the agent works in every period? We denote this payoff by $\overline{W}(\mathbf{1}_q, q)$. The principal's payoff in the corresponding equilibrium poses a lower bound on the principal's equilibrium payoff.

We compare this bound on the principal's equilibrium payoff with the principal's payoff in the no-delay principal-optimum Markov equilibrium. Since the total surplus is fixed by the convention that the agent works in every period, we can do this by comparing the agent's payoff in the two equilibria. In particular, for any q we

- construct an equilibrium in which the agent always works, giving the agent payoff $\overline{W}(\mathbf{1}_q, q)$,
- show $\overline{W}(\mathbf{1}_q, q) \leq W^*(\mathbf{1}_q, q)$,
- show that $\overline{W}(\mathbf{1}_q, q)$ converges to $W^*(\mathbf{1}_q, q)$ as Δ gets small.

This gives us a lower bound on the principal's payoff that is tight (since we have an equilibrium achieving the payoff) and that converges to the Markov payoff as Δ gets small. Notice that $\overline{W}(\mathbf{1}_q, q)$ is also an upper bound on the agent's payoff. The equilibrium

we construct maximizes the surplus and gives the principal her lowest payoff. Any other equilibrium must feature a (weakly) higher payoff for the principal and a (weakly) smaller surplus, and hence can only decrease the agent's payoff.

Let $\tau(\Delta, q)$, typically written simply as τ , be the number of failed experiments required to push the posterior expectation below the threshold \underline{q} for abandoning the project. We then denote the corresponding posteriors by $q_\tau, q_{\tau-1}, \dots, q_1$, with $q_\tau = q$ and with q_1 satisfying

$$\frac{(1-p)q_1}{1-pq_1} < \underline{q} = \frac{2c}{p\pi}.$$

Hence, if an experiment is undertaken at posterior q_1 , no further experimentation will occur.

A.21.1 The No-Delay Principal-Optimum Markov Equilibrium

We start with the no-delay principal-optimum Markov equilibrium. In the last period, we have

$$W^*(\mathbf{1}_{q_1}, q_1) = c.$$

In general, we have

$$\begin{aligned} W^*(\mathbf{1}_{q_\tau}, q_\tau) &= pq_\tau\pi(1-s_\tau) + \delta(1-pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &= c + \delta\frac{q_\tau}{q_{\tau-1}}W^*(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}), \end{aligned}$$

where the second equality is the incentive constraint. Using this second equality to iterate, we have

$$\begin{aligned} W^*(\mathbf{1}_{q_\tau}, q_\tau) &= c + \delta\frac{q_\tau}{q_{\tau-1}}W^*(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &= c + \delta\frac{q_\tau}{q_{\tau-1}}\left(c + \delta\frac{q_{\tau-1}}{q_{\tau-2}}W^*(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2})\right) \\ &= c + \delta\frac{q_\tau}{q_{\tau-1}}c + \delta^2\frac{q_\tau}{q_{\tau-2}}W^*(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2}) \\ &= c + \delta\frac{q_\tau}{q_{\tau-1}}c + \delta^2\frac{q_\tau}{q_{\tau-2}}c + \delta^3\frac{q_\tau}{q_{\tau-3}}W^*(\mathbf{1}_{q_{\tau-3}}, q_{\tau-3}) \\ &\quad \vdots \\ &= cq_\tau\left[\frac{1}{q_\tau} + \frac{\delta}{q_{\tau-1}} + \frac{\delta^2}{q_{\tau-2}} + \frac{\delta^3}{q_{\tau-3}} + \dots + \frac{\delta^{\tau-2}}{q_2}\right] + \delta^{\tau-1}\frac{q_\tau}{q_1}W^*(\mathbf{1}_{q_1}, q_1) \\ &= cq_\tau\left[\frac{1}{q_\tau} + \frac{\delta}{q_{\tau-1}} + \frac{\delta^2}{q_{\tau-2}} + \frac{\delta^3}{q_{\tau-3}} + \dots + \frac{\delta^{\tau-2}}{q_2} + \frac{\delta^{\tau-1}}{q_1}\right]. \end{aligned} \tag{61}$$

A.21.2 The Bound

We now ask what would be the most the principal would have to pay each period, in order to get the agent to work, and what would be the agent's resulting payoff. We proceed recursively. First, we set

$$\overline{W}(\mathbf{1}_{q_1}, q_1) \in [c, pq_1\pi - c].$$

This is simply the statement that in the final period, the payoff of the agent is bounded by the Markov payoff and the entire surplus.

There are two possibilities for the largest amount the principal must pay the agent to work. First, it may be that the agent is indifferent between working and shirking, and any larger value of s induces the agent to shirk. In this case, we can write the agent's incentive constraint as

$$pq\pi(1 - s^W) + \delta(1 - pq)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) = c + \delta \frac{q_\tau}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}),$$

and hence the agents' maximum value as

$$\overline{W}(\mathbf{1}_{q_\tau}, q_\tau) = c + \delta \frac{q_\tau}{q_{\tau-1}} \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}).$$

Alternatively, it may be that the agent's incentive constraint is slack and the agent strictly prefers to work. If the agent strictly prefers to work, why doesn't the principal increase s to some $s + \varepsilon$? The equilibrium presumption is that if the principal does so, the agent shirks, consuming the advance c and prompting no belief revision. For this to be suboptimal, it must be that

$$c + \delta W(\mathbf{1}_{q_\tau}, q_\tau) \geq (1 - (s_\tau + \varepsilon))pq_\tau\pi + \delta(1 - pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}).$$

Since this must hold for every $\varepsilon > 0$, it must hold for the limiting case of $\varepsilon = 0$ (the most stringent form of the inequality). We can also focus on the case in which this constraint holds with equality, since this will fix the bound on s_τ . Hence, the relevant agent's incentive constraint is then

$$c + \delta W(\mathbf{1}_{q_\tau}, q_\tau) = (1 - s_\tau^S) pq_\tau \pi + \delta(1 - pq_\tau) W(\mathbf{1}_{q_\tau}, q_{\tau-1}). \quad (62)$$

We can rearrange (62) to obtain

$$pq_\tau \pi s_\tau^S = pq_\tau \pi - c - \delta w(q_\tau, q_\tau) + \delta(1 - pq) W(\mathbf{1}_{q_\tau}, q_{\tau-1}).$$

How small can we make s_τ^S ? The tools we have for doing this are the continuation payoffs $W(\mathbf{1}_{q_\tau}, q_\tau)$ and $W(\mathbf{1}_{q_\tau}, q_{\tau-1})$. We would like to make the former as large as possible, and

the latter as small as possible. However, these are not independent. A lower bound on the latter is given by the fact that

$$W(\mathbf{1}_{q_\tau}, q_{\tau-1}) \geq \frac{q_{\tau-1}}{q_\tau} W(\mathbf{1}_{q_\tau}, q_\tau),$$

since a pessimistic agent can always duplicate the actions of a more optimistic agent. Hence, no matter what choice we make for $W(\mathbf{1}_{q_\tau}, q_\tau)$, the smallest we can make s_τ is the solution to

$$pq_\tau \pi s_\tau^S = pq_\tau \pi - c - \delta W(\mathbf{1}_{q_\tau}, q_\tau) + \delta(1 - pq) \frac{q_{\tau-1}}{q_\tau} W(\mathbf{1}_{q_\tau}, q_\tau).$$

We can reformulate this condition as

$$pq_\tau \pi s_\tau^S = pq_\tau \pi - c - \delta \left(1 - (1 - pq) \frac{q_{\tau-1}}{q_\tau} \right) W(\mathbf{1}_{q_\tau}, q_\tau).$$

Now it is apparent that we want to make $W(\mathbf{1}_{q_\tau}, q_\tau)$ as large as possible.

We claim that an upper bound on $W(\mathbf{1}_{q_\tau}, q_\tau)$ is $\overline{W}(\mathbf{1}_{q_\tau}, q_\tau)$, the largest bound available to the agent when the agent always works. Suppose we have an alternative candidate equilibrium giving the agent a larger payoff. Then the equilibrium must involve some delay, and hence must involve a smaller surplus than the equilibrium giving payoff $\overline{W}(\mathbf{1}_{q_\tau}, q_\tau)$. It must then involve a smaller payoff from the principal than the payoff the principal receives from the most expensive way of inducing the agent to always work. But since the principal has the option of always inducing the agent to work, the candidate equilibrium cannot be an equilibrium. Instead, the principal would induce work, earning a higher payoff even if this must be done in its most expensive way.

Hence, we have

$$pq_\tau \pi s_\tau^S = pq_\tau \pi - c - \delta \left(1 - (1 - pq) \frac{q_{\tau-1}}{q_\tau} \right) \overline{W}(\mathbf{1}_{q_\tau}, q_\tau).$$

We can then calculate

$$\begin{aligned} \overline{W}(\mathbf{1}_{q_\tau}, q_\tau) &= (1 - s_\tau^S) pq_\tau \pi + \delta(1 - pq) \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &\geq c + \delta \left((1 - (1 - pq) \frac{q_{\tau-1}}{q_\tau}) \overline{W}(\mathbf{1}_{q_\tau}, q_\tau) + \delta(1 - pq) \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \right) \\ &= c + \delta p \overline{W}(\mathbf{1}_{q_\tau}, q_\tau) + \delta(1 - pq) \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &= \frac{1}{1 - \delta p} c + \delta \frac{1 - pq}{1 - \delta p} \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &= \frac{1}{1 - \delta p} c + \delta \frac{1 - p}{1 - \delta p} \frac{q_\tau}{q_{\tau-1}} \overline{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \end{aligned}$$

Combining these calculations, an upper bound on the agent's payoff is given by the solution to the difference equation

$$\bar{W}(\mathbf{1}_{q_\tau}, q_\tau) = \max \left\{ c, \frac{1}{1 - \delta p} c \right\} + \max \left\{ \delta \frac{q_\tau}{q_{\tau-1}}, \delta \frac{1-p}{1 - \delta p} \frac{q_\tau}{q_{\tau-1}} \right\} \bar{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}).$$

Taking the maximum in each case, we can write

$$\begin{aligned} \bar{W}(\mathbf{1}_{q_\tau}, q_\tau) &= \frac{1}{1 - \delta p} c + \delta \frac{q_\tau}{q_{\tau-1}} \bar{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ &:= B + A_\tau \bar{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \end{aligned}$$

We now solve for

$$\begin{aligned} \bar{W}(\mathbf{1}_{q_\tau}, q_\tau) &= A_\tau \bar{W}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) + B \\ &= A_\tau A_{\tau-1} \bar{W}(\mathbf{1}_{q_{\tau-2}}, q_{\tau-2}) + A_\tau B + B \\ &= A_\tau A_{\tau-1} A_{\tau-2} \bar{W}(\mathbf{1}_{q_{\tau-3}}, q_{\tau-3}) + A_\tau A_{\tau-1} B + A_\tau B + B \\ &= A_\tau A_{\tau-1} A_{\tau-2} A_{\tau-3} \bar{W}(\mathbf{1}_{q_{\tau-4}}, q_{\tau-4}) + A_\tau A_{\tau-1} A_{\tau-2} B + A_\tau A_{\tau-1} B + A_\tau B + B \\ &\vdots \\ &= A_\tau \cdots A_2 \bar{W}(\mathbf{1}_{q_1}, q_1) \\ &\quad + A_\tau \cdots A_3 B \\ &\quad + A_\tau \cdots A_4 B \\ &\vdots \\ &\quad + A_\tau A_{\tau-1} B \\ &\quad + A_\tau B \\ &\quad + B. \end{aligned} \tag{63}$$

Now we compare (61) and (63), holding fixed the posterior q that comprises our point of departure, but allowing Δ to approach zero and hence $\tau(\Delta)$ to grow large. The final term in the equilibrium payoff $W^*(\mathbf{1}_{q_{\tau(\Delta)}}, q_{\tau(\Delta)})$ given by (61) is

$$\delta(\Delta)^{\tau(\Delta)-1} \frac{q_{\tau(\Delta)}}{q_1} W^*(\mathbf{1}_{q_1}, q_1) = \delta(\Delta)^{\tau(\Delta)-1} \frac{q_{\tau(\Delta)}}{q_1} c \Delta,$$

while our bound (63) has as its corresponding term

$$A_{\tau(\Delta)} \cdots A_2 \bar{W}(\mathbf{1}_{q_1}, q_1) \leq \delta(\Delta)^{\tau(\Delta)-1} \frac{q_{\tau(\Delta)}}{q_1} \bar{W}(\mathbf{1}_{q_1}, q_1) = \delta(\Delta)^{\tau(\Delta)-1} \frac{q_{\tau(\Delta)}}{q_1} (q_1 p \pi - c) \Delta.$$

We then note that both terms approach zero as does Δ . The sum of the remaining terms comprising $W^*(\mathbf{1}_{q_{\tau(\Delta)}}, q_{\tau(\Delta)})$ in (61) is given by $cq_{\tau(\Delta)}\Delta$ times

$$\frac{1}{q_{\tau(\Delta)}} + \frac{\delta(\Delta)}{q_{\tau(\Delta)-1}} + \frac{\delta(\Delta)^2}{q_{\tau(\Delta)-2}} + \frac{\delta(\Delta)^3}{q_{\tau(\Delta)-3}} + \dots + \frac{\delta(\Delta)^{\tau(\Delta)-2}}{q_2}.$$

Under our bound, the corresponding term in (63) is $cq_{\tau(\Delta)}\Delta$ times

$$\frac{1}{1 - \delta(\Delta)p\Delta} \left[\frac{1}{q_{\tau(\Delta)}} + \frac{\delta(\Delta)}{q_{\tau(\Delta)-1}} + \frac{\delta(\Delta)^2}{q_{\tau(\Delta)-2}} + \frac{\delta(\Delta)^3}{q_{\tau(\Delta)-3}} + \dots + \frac{\delta(\Delta)^{\tau(\Delta)-2}}{q_2} \right].$$

But $\frac{1}{1 - \delta(\Delta)p\Delta} \rightarrow 1$ as time periods get short, while the common term is bounded, and hence $\lim_{\Delta \rightarrow 0} \overline{W}(\mathbf{1}_q, q) \leq W^*(\mathbf{1}_q, q)$, giving the result.

We now show that for an interval $[q, \tilde{q}]$ of priors and for sufficiently small Δ , the upper bound we have calculated on the agent's payoff is tight. The requirement on the interval of priors is that it be such that $s^S < s^W$, which we know is the case for a lower interval that remains nondegenerate as $\Delta \rightarrow 0$.

We used one approximation in the course of constructing the bound on the agent's payoff, namely that

$$W(\mathbf{1}_{q_{\tau}}, q_{\tau-1}) \geq \frac{q_{\tau-1}}{q_{\tau}} W(\mathbf{1}_{q_{\tau}}, q_{\tau}).$$

It thus suffices to show that this is an equality for the range of priors in question. To do this, it suffices to show that a pessimistic agent (one with posterior $q_{\tau-1}$) will work in every period, given that the principal's current belief is q_{τ} , and given the equilibrium that we have constructed, involving offer s^S in every period. Notice that we know a pessimistic agent will not do so when the share offered in every period is s^W . In that case, the pessimistic agent shirks at first opportunity and then has a belief matching the degenerate belief of the principal. However, we are now assuming that share s^S is offered in each period, which is more generous to the agent, making work more attractive.

We argue by induction. Suppose the last period has been reached, meaning that the principal is characterized by a belief q_1 . On the equilibrium path, the principal offers share s^S , which induces the agent to work. We must show that this offer also induces work from an agent characterized by belief $q_0 = \varphi(q_1)$ to work. From the incentive constraint fixing s^S , we have

$$pq_1\pi(1 - s^S) = c + \delta pq_1\pi(1 - s^S) - \delta(1 - pq_1)W(\mathbf{1}_{q_1}, q_0). \quad (64)$$

We need to show

$$pq_0\pi(1 - s^S) \geq c, \quad (65)$$

which suffices for an agent characterized by prior q_0 to work. From (64), we have

$$pq_1\pi(1 - s^S) = \frac{c}{1 - \delta} - \frac{\delta(1 - pq_1)}{1 - \delta} W(\mathbf{1}_{q_1}, q_0).$$

Using this in (65), we need to show

$$\frac{q_0}{q_1} \frac{c}{1-\delta} - \frac{q_0}{q_1} \frac{\delta(1-pq_1)}{1-\delta} \max\{pq_0\pi(1-s^S), c\} \geq c.$$

We suppose that the pessimistic agent shirks and show that this inequality holds, contradicting the supposition that the agent shirks and establishing the result. Taking $\max\{pq_0\pi(1-s^S), c\} = c$, using the updating rule and deleting the common factor c , we have

$$\frac{q_0}{q_1} \frac{1}{1-\delta} - \frac{\delta}{1-\delta} (1-p) \geq 1.$$

A successive series of manipulations gives

$$\begin{aligned} \frac{q_0}{q_1} - \delta(1-p) &\geq 1 - \delta, \\ \frac{q_0}{q_1} &\geq 1 - \delta p, \\ 1-p &\geq (1-\delta p)(1-pq_1), \\ 1-p &\geq 1 - \delta p - pq_1 + \delta p^2 q_1, \\ \delta p + pq_1 &\geq p + \delta p^2 q_1, \\ \delta + q_1 &\geq 1 + p\delta q_1, \end{aligned}$$

which holds for sufficiently small Δ .

Now we turn to the induction step. We consider a belief q_τ and the associated more pessimistic belief $q_{\tau-1} = \varphi(q_\tau)$. The induction hypothesis is that

$$W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-2}) = \frac{q_{\tau-2}}{q_{\tau-1}} W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}).$$

We know, from the definition of s^S , that

$$pq_\tau\pi(1-s^S) = c + \delta W(q_\tau, q_\tau) - \delta(1-pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}).$$

Using the equilibrium definition of $W(q_\tau, q_\tau)$, this is

$$pq_\tau\pi(1-s^S) = c + \delta[pq_\tau\pi(1-s^S) + \delta(1-pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})] - \delta(1-pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}). \quad (66)$$

Our goal is to show that an agent who is one step more pessimistic would prefer to work, or

$$pq_{\tau-1}\pi(1-s^S) + \delta(1-pq_{\tau-1})W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-2}) \geq c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \quad (67)$$

We can reformulate (66) to obtain

$$pq_\tau \pi(1 - s^S) = \frac{c}{1 - \delta} + \frac{\delta}{1 - \delta}(1 - pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \frac{\delta}{1 - \delta}(1 - pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}),$$

and then multiply by $\frac{q_{\tau-1}}{q_\tau}$ and insert in (67) to obtain

$$\begin{aligned} & \frac{q_{\tau-1}}{q_\tau} \frac{c}{1 - \delta} + \frac{q_{\tau-1}}{q_\tau} \frac{\delta}{1 - \delta}(1 - pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \frac{q_{\tau-1}}{q_\tau} \frac{\delta}{1 - \delta}(1 - pq_\tau)W(\mathbf{1}_{q_\tau}, q_{\tau-1}) \\ & \geq c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \delta(1 - pq_{\tau-1})W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-2}). \end{aligned}$$

We use the induction hypothesis to rewrite $W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-2})$ as $\frac{q_{\tau-2}}{q_{\tau-1}}W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})$. It then suffices to assume that the payoff $W(\mathbf{1}_{q_\tau}, q_{\tau-1})$ is generated by a path of play that begins with a shirk, yielding a contradiction that establishes the result. This allows us to rewrite the preceding inequality as

$$\begin{aligned} & \frac{q_{\tau-1}}{q_\tau} \frac{c}{1 - \delta} + \frac{q_{\tau-1}}{q_\tau} \frac{\delta}{1 - \delta}(1 - pq_\tau)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \frac{q_{\tau-1}}{q_\tau} \frac{\delta}{1 - \delta}(1 - pq_\tau)[c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})] \\ & \geq c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \delta(1 - pq_{\tau-1})\frac{q_{\tau-2}}{q_{\tau-1}}W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \end{aligned}$$

Using the updating rules, we can write this as

$$\begin{aligned} & \frac{q_{\tau-1}}{q_\tau} \frac{c}{1 - \delta} + \frac{\delta}{1 - \delta}(1 - p)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \frac{\delta}{1 - \delta}(1 - p)c - \frac{\delta^2}{1 - \delta}(1 - p)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \\ & \geq c + \delta W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) - \delta(1 - p)W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}). \end{aligned}$$

Regrouping terms and applying successive simplifications gives

$$\begin{aligned} c \left[\frac{q_{\tau-1}}{q_\tau} \frac{1}{1 - \delta} - \frac{\delta}{1 - \delta}(1 - p) - 1 \right] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \left[p - \frac{\delta(1 - p)}{1 - \delta} + \frac{\delta^2}{1 - \delta}(1 - p) \right], \\ c \left[\frac{1 - p}{1 - pq_\tau} - \delta(1 - p) - (1 - \delta) \right] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) [p(1 - \delta) - \delta(1 - p) + \delta^2(1 - p)], \\ c [(1 - p) - (1 - \delta p)(1 - pq_\tau)] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) [p - \delta + \delta^2(1 - p)] (1 - pq_\tau), \\ c [1 - p - (1 - \delta p - pq_\tau + \delta p^2 q_\tau)] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) [p - \delta + \delta^2 - \delta^2 p] (1 - pq_\tau), \\ c [-p + \delta p + pq_\tau - \delta p^2 q_\tau] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) [p - \delta + \delta^2 - \delta^2 p] (1 - pq_\tau), \\ c [\delta + q_\tau - 1 - \delta p q_\tau] & \geq W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \left[1 - \delta^2 - \frac{\delta(1 - \delta)}{p} \right] (1 - pq_\tau). \end{aligned}$$

As $\Delta \rightarrow 0$, the coefficient on c on the left side approaches q_τ . On the right side, $(1 - pq_\tau)$ approaches 1, and the coefficient on $W(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})$ approaches $-\frac{r}{p}$, ensuring the inequality.

A.22 Proof of Lemma 16

Let $\underline{V}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1})$ be the smallest principal payoff available from a no-delay Markov path of play, given the common belief $q_{\tau-1}$, and let $\overline{V}(\mathbf{1}_{q_\tau}, q_\tau)$ be the largest such payoff, given the common belief q_τ , with $q_{\tau-1} = \varphi(q_\tau)$. We need to show that $\underline{V}(\mathbf{1}_{q_{\tau-1}}, q_{\tau-1}) \leq \overline{V}(\mathbf{1}_{q_\tau}, q_\tau)$.

Let S_τ be the surplus available at posterior q_τ and $S_{\tau-1}$ be the surplus available at posterior $q_{\tau-1}$. Let, conserving on notation, \underline{W}_τ be the smallest agent payoff from a no-delay path of play, given common belief q_τ , and let $\overline{W}_{\tau-1}$ be the largest agent no-delay payoff given belief $q_{\tau-1}$. Then we need to show

$$S_\tau - \underline{W}_\tau \geq S_{\tau-1} - \overline{W}_{\tau-1}.$$

A sequence of manipulations gives the equivalent statements:

$$\begin{aligned} (pq_\tau\pi - c) + \delta(1 - pq_\tau)S_{\tau-1} &\geq S_{\tau-1} - \overline{W}_{\tau-1} + \underline{W}_\tau, \\ (pq_\tau\pi - c) + \delta(1 - pq_\tau)S_{\tau-1} &\geq S_{\tau-1} - \overline{W}_{\tau-1} + c + \delta \frac{q_\tau}{q_{\tau-1}} \underline{W}_{\tau-1}, \\ \frac{pq_\tau\pi - 2c}{1 - \delta(1 - pq_\tau)} &\geq S_{\tau-1} - \frac{\overline{W}_{\tau-1} - \delta \frac{q_\tau}{q_{\tau-1}} \underline{W}_{\tau-1}}{1 - \delta(1 - pq_\tau)}. \end{aligned}$$

We have $\frac{pq_\tau\pi - 2c}{1 - \delta(1 - pq_\tau)} \geq S_{\tau-1} - \overline{W}_{\tau-1}$, since the left side is the payoff the principal would receive if the principal had to pay only c in each period and if failures did not diminish the posterior, and hence it suffices to show

$$\frac{\overline{W}_{\tau-1} - \delta \frac{q_\tau}{q_{\tau-1}} \underline{W}_{\tau-1}}{1 - \delta(1 - pq_\tau) \overline{W}_{\tau-1}} \leq 1 + \varepsilon,$$

for some $\varepsilon > 0$. Another sequence of manipulations gives

$$\begin{aligned} 1 - \delta \frac{q_\tau}{q_{\tau-1}} \frac{\underline{W}_{\tau-1}}{\overline{W}_{\tau-1}} &\leq (1 - \delta(1 - pq_\tau))(1 + \varepsilon), \\ 1 - \delta \frac{q_\tau}{q_{\tau-1}} \frac{\underline{W}_{\tau-1}}{\overline{W}_{\tau-1}} &\leq 1 - \delta(1 - pq_\tau) + \varepsilon(1 - \delta(1 - pq_\tau)), \\ \delta \frac{1 - pq_\tau}{1 - p} \frac{\underline{W}_{\tau-1}}{\overline{W}_{\tau-1}} &\geq \delta(1 - pq_\tau) - \varepsilon(1 - \delta(1 - pq_\tau)), \\ \frac{\underline{W}_{\tau-1}}{(1 - p)\overline{W}_{\tau-1}} &\geq 1 - \varepsilon \frac{1 - \delta(1 - pq_\tau)}{\delta(1 - pq_\tau)}, \\ \frac{\underline{W}_{\tau-1}}{\overline{W}_{\tau-1}} &\geq (1 - p) - \varepsilon(1 - p) \frac{1 - \delta(1 - pq_\tau)}{\delta(1 - pq_\tau)}. \end{aligned}$$

In the course of proving Lemma 15, we have derived an expression for $\underline{W}_{\tau-1}$ and an upper bound on $\overline{W}_{\tau-1}$, and we can insert these to obtain

$$\frac{k_1\Delta + Z}{k_2\Delta + \frac{Z}{1-\delta p}} \geq (1-p) - \varepsilon \frac{1 - \delta(1 - pq_\tau)}{\delta(1 - pq_\tau)},$$

where we will later need that $k_1 = \theta c$ and $k_2 = \theta(pq_1\pi - c)$, and we need to know about θ and Z only that if we fix a posterior q and let Δ get small (holding q constant, so that the number of revisions following q grows), θ and Z are bounded away from zero. Every term r and p should also be multiplied by Δ , but we omit these to improve readability. Now we manipulate to give

$$\begin{aligned} \frac{(1-\delta p)(k_1\Delta + Z)}{(1-\delta p)k_2\Delta + Z} &\geq (1-p) - \varepsilon(1-p) \frac{1 - \delta(1 - pq_\tau)}{\delta(1 - pq_\tau)}, \\ \delta(1-pq)(1-\delta p)(k_1\Delta + Z) &\geq (1-p)(1-pq)\delta[(1-\delta p)k_2\Delta + Z], \\ &\quad - \varepsilon(1-p)(1-\delta(1-pq))[(1-\delta p)k_2\Delta + Z], \\ (1-r)(1-pq)(1-(1-r)p)(k_1\Delta + Z) &\geq (1-p)(1-pq)(1-r)[(1-(1-r)p)k_2\Delta + Z], \\ &\quad - \varepsilon(1-p)(1-(1-r)(1-pq))[(1-(1-r)p)k_2\Delta + Z], \\ (1-r)(1-pq)(1-p+pr)(k_1\Delta + Z) &\geq (1-p)(1-pq)(1-r)[(1-p+pr)k_2\Delta + Z], \\ &\quad - \varepsilon(1-p)(r+pq-rpq)[(1-p+pr)k_2\Delta + Z]. \end{aligned}$$

To evaluate this inequality, we first examine the terms involving Δ^0 , which give us simply $Z \geq Z$, which obviously holds with equality. Hence, we examine terms involving Δ^1 , finding

$$-r\Delta - pqZ - pZ + k_1\Delta \geq -pZ - pqZ - rZ + k_2\Delta - \varepsilon rZ - \varepsilon pqZ.$$

Hence, we need to show (using the definitions of k_1 and k_2 to obtain the second inequality)

$$\begin{aligned} k_1\Delta &\geq k_2\Delta - \varepsilon Z(r + pq), \\ \theta c &\geq \theta(pq_1\pi - c) - \varepsilon Z(r + pq), \\ \varepsilon Z(r + pq) &\geq \theta(pq_1\pi - 2c), \\ \varepsilon Z(r + pq) &\geq 2\theta c \left(\frac{q_1}{q} - 1 \right), \\ \varepsilon Z \frac{(r + pq)}{c} &\geq 2\theta \left(\frac{q_1}{q} - 1 \right). \end{aligned}$$

We now note that the left side is constant in Δ , while the right side approaches zero as does Δ , giving the result.

A.23 Proof of Lemma 18

For the case of $\psi > 2$ and $\psi > \sigma$, in which there is no delay, this result already follows from Lemma 15. Lemma 15 also ensures this is the case for beliefs $q < q^*$ when $\psi > 2$ and $\psi < \sigma$ (delay for high beliefs). It is a straightforward adaptation of Lemma 2, requiring only a substitution of the appropriate initial conditions, to show that this is the case for $q > q^{**}$ when $\psi < 2$ and $\psi > \sigma$ (delay for low beliefs). We thus need to consider periods of delay.

We begin with the following preliminary result. Suppose we have an equilibrium and a period τ with values $v_{\tau-1} > 0$ and $v_\tau = 0$. We bound the amount of delay we can introduce between periods τ_1 and τ . We fix the continuation behavior prescribed by this equilibrium, and then introduce delay between periods $\tau + 1$ and τ , with the total discounting being these periods given by $\Lambda\delta(\Delta)$. We show that as $\Delta \rightarrow 0$, Λ approaches 1.

Suppose that Δ time has passed since the offer was made that caused the belief to be revised from $q_{\tau+1}$ to q_τ . The principal is supposed to wait an additional period of time equivalent to discount factor Λ , and we have an equilibrium only if the principal does not find it profitable to “jump” this waiting time.

We consider two cases. Suppose first that $s_\tau^W \geq s_\tau^S$. We need to formulate the incentive constraint for q_τ . If $s_\tau = s_\tau^W$, as would be the case if the continuation equilibrium were the no-delay principal-optimum equilibrium, the incentive constraint is

$$pq_\tau\pi(1 - s_\tau) + \delta_\tau(1 - pq_\tau)w_{\tau-1} = c + \delta_\tau \frac{q_\tau}{q_{\tau-1}} w_{\tau-1}.$$

However, we have chosen the continuation equilibrium to be such that $v_\tau = 0$, which may not be the no-delay principal-optimum equilibrium, and may call for $s_\tau \in [s_\tau^S, s_\tau^W)$. Then there is some ε such that

$$pq_\tau\pi(1 - s_\tau) + \delta_\tau(1 - pq_\tau)w_{\tau-1} = c + \delta_\tau \frac{q_\tau}{q_{\tau-1}} w_{\tau-1} + \varepsilon,$$

where ε is nonnegative and bounded (for example, by $pq_\tau\pi - 2c$).

Consider what happens if the principal makes an offer $s_\tau + \varepsilon$. The equilibrium calls for the agent to reject this offer, conditional on being expected to reject the offer. If the agent were to accept the offer, it would be profitable for the principal to make it, since the Markov assumption would then force the continuation of the equilibrium play appropriate for belief $q_{\tau-1}$, and the principal would have reached this continuation play more quickly and via a slightly more lucrative offer than the equilibrium prescription. We must allow ε to be arbitrarily small, and hence must show that the agent must find it at least weakly profitable to reject s_τ if made, given that such a rejection is expected. Hence, it must be that

$$c + \Lambda w_\tau \geq pq_\tau\pi(1 - s_\tau) + \Lambda(1 - pq_\tau)[c + \delta w_{\tau-1}].$$

Using the incentive constraint in the latter, this is

$$c + \Lambda w_\tau \geq c + \delta_\tau \frac{q_\tau}{q_{\tau-1}} w_{\tau-1} - \delta_\tau (1 - pq_\tau) w_{\tau-1} + \Lambda (1 - pq_\tau) c + \Lambda (1 - pq_\tau) \delta_\tau w_{\tau-1} + \varepsilon.$$

Eliminating c from each side, substituting for w_τ , and rearranging, this is

$$\Lambda c + \delta_\tau \Lambda \theta_\tau w_{\tau-1} + \Lambda \varepsilon \geq \delta_\tau \theta_\tau w_{\tau-1} - \delta_\tau (1 - p) \theta_\tau w_{\tau-1} + \Lambda (1 - p) \theta_\tau c + \delta_\tau \Lambda (1 - p) \theta_\tau w_{\tau-1},$$

where $\theta_\tau := \frac{q_\tau}{q_{\tau-1}}$. A series of successive rearrangements now gives

$$\begin{aligned} \Lambda c [1 - (1 - p) \theta_\tau] &\geq \theta_\tau w_{\tau-1} [\delta_\tau - \delta_\tau (1 - p) + \delta_\tau \Lambda (1 - p) - \delta_\tau \Lambda] + (1 - \Lambda) \varepsilon, \\ \Lambda c p q_\tau &\geq \theta_\tau w_{\tau-1} (\delta_\tau p - \delta_\tau p \Lambda) + (1 - \Lambda) \varepsilon, \\ \Lambda c q_\tau &\geq \theta_\tau w_{\tau-1} \delta_\tau (1 - \Lambda) + (1 - \Lambda) \frac{\varepsilon}{p}, \\ \frac{\Lambda}{1 - \Lambda} &\geq \frac{\delta_\tau \theta_\tau w_{\tau-1}}{c q_\tau} + \frac{\varepsilon}{p d q_\tau}. \end{aligned}$$

As $\Delta \rightarrow 0$, as long as $\delta_\tau(\Delta)$ does not approach zero, the first term on the right side grows without bound, while the second remains nonnegative. This ensures that Λ converges to 1 as $\Delta \rightarrow 0$.

For the second case, suppose that $s_\tau^W \leq s_\tau^S$. Here, we must ensure that the agent at least weakly prefers to reject an offer s_τ'' , conditional on being expected to reject. It is then immediate that $\Lambda \geq \delta(\Delta)$. The value s_τ'' is by definition a value that makes the agent just indifferent between accepting and rejecting, given that a rejection is expected and that there is delay $\delta(\Delta)$ until the next offer. Should the equilibrium strategies, after the offer that prompted the belief reduction from $q_{\tau+1}$ to q_τ and after a waiting time of length Δ , prescribe further discounting of length exceeding $\delta(\Delta)$, then the agent will strictly prefer to accept offer s_τ'' immediately, which would be profitable for the principal and hence would disrupt the equilibrium.

This in turn allows us to show the following. Fix a posterior q and consider an equilibrium in which $v(q) = 0$. Fix $\varepsilon > 0$ and suppose that continuing with the maximal no-delay program backward from q gives $v_{q+\varepsilon} < 0$. Then for sufficiently small Δ , there exists a $q' \in (q, q + \varepsilon]$ with $v_{q'} = 0$. To show this, number periods so that q occurs at period 0 and $q + \varepsilon$ at period T . We will be interested in the case in which Δ gets small, and so T will depend on Δ . The value v at the posterior $\varphi(q)$ will either be positive, in which case there is no delay between q and $\varphi(q)$ and hence $\delta_0 = \delta(\Delta)$, or the value v at posterior $\varphi(q)$ will equal zero, in which case δ_0 is set by the need to set $v(q) = 0$. In either case, δ_0 will remain bounded as $\Delta \rightarrow 0$.

Suppose the claim fails and hence the principal's payoff is positive over the interval $[q, q + \varepsilon]$. Then over this interval, there can be delay only in period 1, and hence we can

write

$$\begin{aligned}
w_T &= c \\
&+ \delta_T \theta_T c \\
&+ \delta_T \delta_{T-1} \theta_T \theta_{T-1} c \\
&+ \delta_T \delta_{T-1} \delta_{T-2} \theta_T \theta_{T-1} \theta_{T-2} c \\
&\vdots \\
&+ \delta_T \delta_{T-1} \delta_{T-2} \cdots \delta_1 \Lambda_1 \theta_T \theta_{T-1} \theta_{T-2} \cdots \theta_1 c \\
&= c \\
&+ \delta(\Delta) \theta_T c \\
&+ (\delta(\Delta))^2 \theta_T \theta_{T-1} c \\
&+ (\delta(\Delta))^3 \theta_T \theta_{T-1} \theta_{T-2} c \\
&\vdots \\
&+ (\delta(\Delta))^{T-1} \Lambda_1 \theta_T \theta_{T-1} \theta_{T-2} \cdots \theta_1 c.
\end{aligned}$$

But as Δ gets small, $\Lambda_1 \rightarrow 1$, and this agent payoff approaches the agent's payoff under the maximal full-effort construction. The latter payoff ensures that the principal earns a negative payoff at $q + \varepsilon$, a contradiction.

Now consider an interval of beliefs $[q', q'']$ over which the canonical Markov equilibrium features a zero principal payoff. The preceding result ensures that for any Markov equilibrium, the set of posteriors at which the principal receives a zero payoff becomes dense in $[q', q'']$ as Δ gets small. A continuity argument then ensures that all equilibrium payoffs converge to the limiting payoffs.

B Online Appendix: Comparative Statics, The First-Best Policy

The optimal stopping time T solves $q_T = c/p\pi$. Up to second-order terms, we have

$$q_{t+\Delta} = \frac{q_t(1 - p\Delta)}{1 - q_t p\Delta},$$

and hence, the posterior belief that the project is good evolves according to

$$\dot{q}_t = -pq_t(1 - q_t),$$

with $q_0 = \bar{q}$. Integrating, the principal's belief at time $t \leq T$ is then

$$q_t = \left(1 + \frac{1 - \bar{q}}{\bar{q}} e^{pt}\right)^{-1},$$

Therefore, inserting in $q_T = c/p\pi$ and solving, the stopping time is given by

$$T = \frac{1}{p} \left[\ln \left(\frac{p\pi}{c} - 1 \right) + \ln \frac{\bar{q}}{1 - \bar{q}} \right].$$

It is immediate from this expression that the first-best policy operates the project longer when the prior probability \bar{q} is larger, when (holding p fixed) the benefit-cost ratio $p\pi/c$ is larger, and when (holding $p\pi/c$ fixed) the success probability p is smaller.

C Online Appendix: Proofs, Unobservable Effort

C.1 Proof of Lemma 3

C.1.1 The Agent's Highest No-Delay Payoff

To find the agent's lowest equilibrium payoff, we first need to solve explicitly for the agent's highest payoff w^A across Markov equilibria without delay. Let $\{q_\tau\}_{\tau-1}^\infty$ be the sequence of posteriors through which equilibrium beliefs will pass, with $q_{\tau-1} = \varphi(q_\tau)$. Section A.21.2 establishes that the agent's highest payoff satisfies the recursion

$$w^A(q_\tau) = \frac{1}{1 - \delta p} c + \delta \frac{1 - p}{1 - \delta p} \frac{q_\tau}{q_{\tau-1}} w^A(q_{\tau-1}),$$

for all values of q that are below some value bounded above \underline{q} (uniformly in Δ). Let us restrict attention for now to such values. In addition, in the last period (period 1), all the surplus goes to the agent.

The solution to the sequence of beliefs q_τ is given by (39). It is more convenient to work with the value normalized by the belief, and so we define $\omega_\tau^A := w_\tau^A / (q_\tau c)$ (where $w_\tau^A := w^A(q_\tau)$). This gives

$$\omega_{\tau+1}^A = \frac{1 + Q_1 \beta^{\tau+1}}{1 - \delta + \delta \beta} + \frac{\delta \beta}{1 - \delta + \delta \beta} \omega_\tau^A,$$

with $\omega_1^A = B + 1 - Q_1 \beta$, where $B := 2(1 - \underline{q})/\underline{q}$, $\beta = 1 - p$, and $Q_1 := (1 - q_1)/q_1$ is the inverse likelihood ratio in the last period. Because it turns out to be irrelevant for the limits, and simplifies expressions, we set $q_1 = \underline{q}$. Manipulating the difference equation gives

$$\begin{aligned} \omega_\tau^A &= \beta^{\tau-1} \delta^{\tau-1} \frac{((\beta - 2)\delta + 1)(\delta - B(1 - \delta)) + \beta(1 - \delta)Q_1((\beta - 2)\delta + 2)}{(\delta - 1)((\beta - 2)\delta + 1)((\beta - 1)\delta + 1)^{\tau-1}} \\ &\quad + \frac{(1 - \beta)\delta - (1 - \delta) - Q_1(1 - \delta)\beta^\tau}{(1 - \delta)((1 - \beta)\delta - (1 - \delta))}. \end{aligned}$$

C.1.2 The Principal's Lowest Payoff

The principal's lowest payoff \underline{v} (corresponding to the agent's highest payoff) across all Markov equilibria without delay, is given by the difference between total surplus and the agent's highest payoff. Given that we have already solved for the agent's payoff, it is more convenient to compute the surplus. Total surplus (normalized by c and q_τ) satisfies the recursion

$$\mathfrak{s}_{\tau+1} = \frac{p\pi}{c} - \frac{1}{q_{\tau+1}} + \delta(1-p)\mathfrak{s}_\tau = B + 1 - Q_1\beta^{n+1} + \delta\beta\mathfrak{s}_\tau,$$

with $\mathfrak{s}_1 = \omega_1^A$, as the agent gets all the surplus in the last period. This gives

$$\mathfrak{s}_\tau = \frac{\beta^\tau [Q_1(\beta\delta - 1) - \delta^\tau(Q_1(\beta - 2)\delta + Q_1 - \delta + 1)] + (2Q_1 + 1)(1 - \delta)}{(1 - \delta)(1 - \beta\delta)}.$$

Section A.21.2 establishes that the principal's payoff in any equilibrium, Markov or not, must be at least the resulting payoff \underline{v} .

C.1.3 The Agent's Lowest Payoff

We now turn to the lowest payoff of the agent. Let us write \underline{w} for what we will call the *interim* value of the agent's payoff, that is, the agent's payoff given that the mandatory waiting time Δ since the last offer has passed, but before any additional, discretionary delay (if any) has occurred. This discretionary delay cannot drive the principal's payoff below \underline{v} . Accordingly, we peg the principal's interim value to \underline{v} . Working with the *interim* values gives a lower bound to the principal's ex post payoff, i.e., her payoff after delay, so the principal also gets at least her lower bound at the point where she makes the offer (and receives a higher payoff when on the verge of making an offer if there is additional delay and $\underline{v} > 0$). This lower bound will be tight, since the principal's payoff is \underline{v} after the mandatory waiting period Δ since the last offer has passed, and the principal has reached the first point at which she can act.

We capture the possibility of discretionary delay by introducing a variable $\Lambda(q) \leq 1$, representing the additional discounting caused by such delay. We have $\Lambda(q) = 1$ if there is no discretionary delay, and otherwise $\Lambda(q) < 1$, in order to capture the reduction in payoffs imposed by the delay until the offer is made. We then have, by definition of *interim* values,

$$\underline{w}(q) = \Lambda(q) \left(pq\pi(1 - s(q)) + \delta(1 - p)\frac{q}{\tilde{q}}\underline{w}(\tilde{q}) \right) = \Lambda(q) \left(c + \delta\frac{q}{\tilde{q}}\underline{w}(\tilde{q}) \right),$$

where $\tilde{q} = \varphi(q)$ is the posterior belief, as well as

$$\underline{v}(q) = \Lambda(q) \left(pq(q)\pi s - c + \delta(1 - p)\frac{q}{\tilde{q}}\underline{v}(\tilde{q}) \right).$$

Combining,

$$\underline{v}(q)/\Lambda(q) = pq\pi - 2c + \delta \frac{q}{\tilde{q}} [(1-p)\underline{v}(\tilde{q}) - p\underline{w}(\tilde{q})],$$

which can be solved for $\Lambda(q)$. Hence, plugging back into the recursion for \underline{w} ,

$$\underline{w}(q) = \frac{v(q)}{pq\pi - 2c + \delta \frac{q}{\tilde{q}} [(1-p)v(\tilde{q}) - p\underline{w}(\tilde{q})]} \left(c + \delta \frac{q}{\tilde{q}} \underline{w}(\tilde{q}) \right).$$

This is a discrete-time Riccati equation that converges pointwise to the continuous-time Riccati equation (19) of Section 3.2.1. The same Riccati equation obtains if we work with the principal's best Markov equilibrium without delay: the choice is irrelevant to the evolution of this lowest payoff, but it is key to the boundary condition that determines the solution to this Riccati equation. Because this boundary condition is at \underline{q} , our restriction to low beliefs (necessary to assert that the lower bound on the principal's equilibrium payoffs, w^A , is actually an equilibrium payoff, as shown in Section A.21.2), is innocuous.

To complete the argument, we must show that the boundary condition at \underline{q} selects the lower of the two solutions identified Section 3.2.1. This requires a fine analysis of the game with $\Delta > 0$. To this end, let us define $\underline{\omega}(q) := \underline{w}$ and $(q)/(cq), \underline{\nu}(q) := \underline{v}(q)/(cq)$ (and also $\underline{\nu}_\tau, \underline{\omega}_\tau$ for $q = q_\tau$). Rearranging the Riccati equation, we get

$$\underline{\omega}_{\tau+1}\underline{\omega}_\tau - \frac{p\pi/c - 2/q_{\tau+1} + \delta(1-p)\underline{\nu}_\tau}{\delta p} \underline{\omega}_{\tau+1} + \frac{\underline{\nu}_{\tau+1}}{p} \underline{\omega}_\tau + \frac{1}{\delta p q_{\tau+1}} \underline{\nu}_{\tau+1} = 0. \quad (68)$$

We now must insert our explicit solution for $\underline{\nu}_\tau$, given our formulas for s_τ and ω_τ^A , namely, $\underline{\nu}_\tau = s_\tau - \omega_\tau^A$. Further, let $a_\tau^\Delta := \underline{\omega}_\tau \Delta / (q_\tau - q_1)$, $\tau > 1$. Note that, given τ , this is a measure of the slope of $\underline{\omega}$ at $q_1 \rightarrow \underline{q}$. Inserting (a^Δ) , we get that $a_{\tau+1}^\Delta$ solves

$$\begin{aligned} a_{\tau+1}^\Delta a_\tau^\Delta - \frac{p\pi/c - 2/q_{\tau+1} + \delta(1-p)(s_\tau - \omega_\tau^A)\Delta}{\delta p(q_\tau - q_1)} a_{\tau+1}^\Delta \\ + \frac{(s_{\tau+1} - \omega_{\tau+1}^A)\Delta}{p(q_{\tau+1} - q_1)} a_\tau^\Delta + \frac{(s_{\tau+1} - \omega_{\tau+1}^A)\Delta}{\delta p q_{\tau+1} (q_n - q_1)(q_{\tau+1} - q_1)} = 0. \end{aligned}$$

Taking limits as $\Delta \rightarrow 0$, this gives, for all τ ,

$$\lim_{\Delta \rightarrow 0} a_\tau^\Delta = \frac{\tau(3+\tau)}{(\tau+1)^2} \frac{(Q_1 - 1)(Q + l_1)^3}{4Q_1^2 r \sigma},$$

and so

$$\lim_{\tau \rightarrow \infty} \lim_{\Delta \rightarrow 0} a_\tau^\Delta = a_\infty := \frac{(Q_1 - 1)(Q_1 + 1)^3}{4Q_1^2 r \sigma}.$$

Given the definition of $\underline{\omega}_\tau$, this implies that $\underline{w}_\tau \Delta / (q_\tau - q_1)$ converges precisely to $\underline{w}'(\underline{q})$, as given in Section 3.2.1, which was to be shown. For reference, if we compute the same limit for $a_\tau^{A,\Delta} := \lim_{\Delta \rightarrow 0} w_\tau^A \Delta / (q_\tau - q_1)$, we get

$$\lim_{\tau \rightarrow \infty} \lim_{\Delta \rightarrow 0} a_\tau^{\Delta,A} = a_\infty^A := \frac{(Q_1 + 1)^3}{Q_1 r \sigma},$$

and we note that $a_\infty < a_\infty^A$, i.e. we have created some “slack” between our new lower bound and the Markov equilibrium payoff (note that $Q_1 > 2$ so $Q_1 - 1 > 0$). One can check that the same limit obtains with the Markov payoff (independently of the initial condition ω_1, ν_1), i.e. $a_\infty^M = \lim_{\tau \rightarrow 0} \lim_{\Delta \rightarrow 0} \omega_\tau \Delta / (q_\tau - q_1) = a_\infty^A$: that is, the averages of the agent’s canonical (or highest) Markov payoff select the solution w_M to the Riccati equation in the continuous-time limit, as should be expected.

C.2 Proof of Proposition 2

Given $w(q)$, we would like to characterize the value of $v(q)$ that maximizes the principal’s payoff among the set of equilibrium payoffs that give the agent $w(q)$. We simplify the argument by assuming the players have access to a public randomization device, describing at the end of the proof the modifications required if no such device is available. We make critical use of the fact that we have a “worst equilibrium” that simultaneously delivers the worst possible equilibrium payoffs to both the principal and the agent.

Fix a value q . Let $(\underline{w}(q), \underline{v}(q))$ denote the worst equilibrium payoffs given q . Let $\tilde{\mathcal{E}}$ be the set of payoffs corresponding to equilibria in which there is no initial delay, i.e., in which the principal immediately makes a serious offer. Then the set of equilibrium payoffs $\mathcal{E}(q)$ is given by that part of the convex hull of $(0, 0)$ and $\tilde{\mathcal{E}}(q)$ in which the principal receives payoff at least $\underline{v}(q)$.³⁴

Fix an equilibrium $w(q)$ and let $\tilde{v}(q)$ be the largest principal payoff consistent with an equilibrium in which the agent receives $w(q)$. We argue that payoffs $(w(q), \tilde{v}(q))$ can be achieved as a convex combination of the worst equilibrium and an equilibrium in $\tilde{\mathcal{E}}$. If $(w(q), \tilde{v}(q))$ is itself in $\tilde{\mathcal{E}}$, the result is immediate, as it is if $(w(q), \tilde{v}(q))$ is not contained in $\tilde{\mathcal{E}}$ but the line passing through $(\underline{w}(q), \underline{v}(q))$ and $(w(q), \tilde{v}(q))$ intersects $\tilde{\mathcal{E}}$.³⁵ Now suppose

³⁴Any equilibrium outcome must consist of a (possibly null) initial delay followed by an offer. The resulting payoff is then a convex combination of the continuation payoff following the offer, which lies in $\tilde{\mathcal{E}}$, and $(0, 0)$. To be an equilibrium outcome, the principal must obviously receive a payoff of at least $\underline{v}(q)$. In addition, the principal must find it optimal to endure the initial delay rather than make an earlier offer. As long as the equilibrium payoff is at least $\underline{v}(q)$, we can construct an equilibrium in which an earlier offer triggers an immediate switch to a continuation equilibrium with payoff $\underline{v}(q)$, giving the result.

³⁵Notice that $(\underline{w}(q), \underline{v}(q))$ is weakly smaller than any element of \mathcal{E} , and since \mathcal{E} is a collection of convex combinations of $(0, 0)$ and points in the convex set $\tilde{\mathcal{E}}$, the payoff $(w(q), \tilde{v}(q))$ must lie between $(\underline{w}(q), \underline{v}(q))$ and some point in $\tilde{\mathcal{E}}$, on the line containing the latter.

that the line passing through $(\underline{w}(q), \underline{v}(q))$ and $(w(q), \tilde{v}(q))$ does not intersect $\tilde{\mathcal{E}}$. The point $(w(q), \tilde{v}(q))$ is itself a convex combination of $(0, 0)$ and a point $(w'(q), v'(q))$ in $\tilde{\mathcal{E}}$. If $(w'(q), v'(q))$ lies above the (non-negatively sloped) line passing through $(\underline{w}(q), \underline{v}(q))$ and $(w(q), \tilde{v}(q))$, then there is a convex combination of $(\underline{w}(q), \underline{v}(q))$ and $(w'(q), v'(q))$ that gives the agent payoff $w(q)$ and the principal a payoff exceeding $\tilde{v}(q)$, a contradiction to the hypothesis that $\tilde{v}(q)$ is the highest principal payoff consistent with an equilibrium in which the agent receives $w(q)$. If $(w'(q), v'(q))$ lies below the line passing through $(\underline{w}(q), \underline{v}(q))$ and $(w(q), \tilde{v}(q))$ (or to the right, if the latter is vertical), then the convex combination of $(0, 0)$ and $(w'(q), v'(q))$ that gives the principal payoff $\underline{v}(q)$ gives the agent a payoff smaller than $\underline{w}(q)$, a contradiction to the hypothesis that $(\underline{w}(q), \underline{v}(q))$ is the worst equilibrium payoff.

Any equilibrium payoff can thus be generated as a mixture between the worst equilibrium and a continuation equilibrium with no delay. (In many cases, the weight on the latter equilibrium will be one.) Hence, we can restrict attention to equilibria generated by sequences $(x(q), s(q))$, where, given posterior q , the worst equilibrium (given the current posterior) is played with probability $1 - x(q)$ (determined by a public randomization device); and if not, a share $s(q)$ is offered in that period that induces the agent to work.

C.2.1 A Preliminary Inequality

We fix a posterior probability and let $w(q)$ and $v(q)$ be equilibrium values, with $\underline{w}(q)$ and $\underline{v}(q)$ being the values of the worst equilibrium given that posterior. Now, let ζ be such that for *any* such posterior probability,

$$\frac{v(q) - \underline{v}(q)}{w(q) - \underline{w}(q)} \leq \zeta.$$

Our first step is to place an upper bound on ζ .

We first argue that there exists a bound on $\frac{v(q) - \underline{v}(q)}{w(q) - \underline{w}(q)}$. Let $\mathfrak{s}(q) = w(q) + v(q)$ be the total surplus generated to the equilibrium. To begin, we consider the case in which the worst equilibrium payoff is $(0, 0)$. We can then show that

$$\frac{\mathfrak{s}(q)}{w(q)} \leq \frac{p\pi - c}{c},$$

which suffices to bound $v(q)/w(q)$. The bound on $\mathfrak{s}(q)/w(q)$ follows from noting that surplus is generated only if the agent exerts effort. Any exertion of effort on the part of the agent creates a discounted surplus of $(p\pi - c)\Delta$, where the discounting reflects the delay until the effort is exerted and the probability that the interaction may terminate before reaching such effort. Of this surplus, at least $c\Delta$ must go to the agent, since otherwise the agent's incentive constraint is surely violated. The ratio $\mathfrak{s}(q)/w(q)$ can then never exceed $(p\pi - c)/c$.

Now suppose the worst equilibrium payoff is not $(0, 0)$. The preceding argument establishes that $\underline{v}(q)/\underline{w}(q)$ is bounded. It then suffices to show that

$$\frac{v(q) - \underline{v}(q)}{w(q) - \underline{w}(q)} \leq \frac{\underline{v}(q)}{\underline{w}(q)}.$$

Suppose this is not the case. Then the line connecting $(w(q), v(q))$ with $(0, 0)$ passes above $(\underline{w}(q), \underline{v}(q))$. This ensures that there exists a convex combination of $(w(q), v(q))$ and $(0, 0)$ that provides payoff $\underline{v}(q)$ to the principal and a payoff lower than $\underline{w}(q)$ to the agent, contradicting the hypothesis that $(\underline{w}(q), \underline{v}(q))$ is the lowest equilibrium payoff.

We now seek an upper bound on the bound ζ . Fix a posterior q . We first note that

$$\begin{aligned} v(q) &= x(q) [(pqs\pi - c)\Delta + \delta(\Delta)(1 - pq\Delta)[x(\varphi(q))v(\varphi(q)) + (1 - x(\varphi(q)))\underline{v}(\varphi(q))]] \\ &\quad + (1 - x(q))\underline{v}(q), \\ w(q) &= x(q) [pq(1 - s)\pi\Delta + \delta(\Delta)(1 - pq\Delta)[x(\varphi(q))w(\varphi(q)) + (1 - x(\varphi(q)))\underline{w}(\varphi(q))]] \\ &\quad + (1 - x(q))\underline{w}(q) \\ &\geq x(q) [c\Delta + \delta(\Delta)[x(\varphi(q))\theta w(\varphi(q)) + (1 - x(\varphi(q)))\theta \underline{w}(\varphi(q))]] + (1 - x(q))\underline{w}(q), \end{aligned}$$

where $\varphi(q)$ is the posterior belief obtained from q given a failure (cf. (31)). The inequality is the agent's incentive constraint and $\theta > 1$ is given by

$$\theta = \frac{q}{\varphi(q)} = \frac{1 - p\Delta q}{1 - p\Delta},$$

and hence is the ratio of the current posterior to next period's posterior, given a failure. We have used here the fact that the continuation values, relevant for posterior $\varphi(q)$, can be written as convex combinations of equilibrium payoffs $(w(\varphi(q)), v(\varphi(q)))$ and the worst equilibrium payoffs $(\underline{w}(\varphi(q)), \underline{v}(\varphi(q)))$. Note that in writing this convex combination, we take $w(\varphi(q))$ and $v(\varphi(q))$ to be the interim values, i.e., values at the point at which the posterior is $\varphi(q)$ and precisely Δ time has elapsed since the previous offer. The equilibrium generating these values may yet entail some delay.

Let us simplify the notation by letting $x(q) = x$, $v(q) = v$, $w(q) = w$, $\underline{v}(q) = \underline{v}$, $\underline{w}(q) = \underline{w}$, $x(\varphi(q)) = \tilde{x}$, $v(\varphi(q)) = \tilde{v}$, $w(\varphi(q)) = \tilde{w}$, $\underline{v}(\varphi(q)) = \underline{\tilde{v}}$, and $\underline{w}(\varphi(q)) = \underline{\tilde{w}}$, and let us drop the explicit representation of Δ . Setting an equality in the agent's incentive constraint and rearranging gives

$$pqs\pi = (pq\pi - c) + \delta(1 - pq)[\tilde{x}\tilde{w} + (1 - \tilde{x})\underline{\tilde{w}}] - \delta[\tilde{x}\theta\tilde{w} + (1 - \tilde{x})\theta\underline{\tilde{w}}].$$

Using this to eliminate the variable s from the value functions gives

$$\begin{aligned} v - \underline{v} &= x [(pq\pi - 2c) + \delta(1 - pq)[\tilde{x}\tilde{w} + (1 - \tilde{x})\underline{\tilde{w}}] - \delta[\tilde{x}\theta\tilde{w} + (1 - \tilde{x})\theta\underline{\tilde{w}}]] \\ &\quad + \delta(1 - pq)[\tilde{x}\tilde{v} + (1 - \tilde{x})\underline{\tilde{v}}] - \underline{v}, \end{aligned} \tag{69}$$

$$w - \underline{w} = x [c + \delta[\tilde{x}\theta\tilde{w} + (1 - \tilde{x})\theta\underline{\tilde{w}}] - \underline{w}]. \tag{70}$$

Dividing (69) by (70), we obtain

$$\frac{v - \underline{v}}{w - \underline{w}} = \frac{(pq\pi - 2c) + [\delta(1 - pq) - \delta\theta]\tilde{x}[\tilde{w} - \underline{\tilde{w}}] + [\delta(1 - pq) - \delta\theta]\underline{\tilde{w}}}{c + \delta\theta[\tilde{x}(\tilde{w} - \underline{\tilde{w}}) + \underline{\tilde{w}}] - \underline{w}} + \frac{\delta(1 - pq)\tilde{x}(\tilde{v} - \underline{\tilde{v}}) + \delta(1 - pq)[\underline{\tilde{v}} - \underline{v}]}{c + \delta\theta[\tilde{x}(\tilde{w} - \underline{\tilde{w}}) + \underline{\tilde{w}}] - \underline{w}}.$$

Using the fact that $\tilde{v} - \underline{\tilde{v}} \leq \zeta(\tilde{w} - \underline{\tilde{w}})$, we can substitute and rearrange to obtain an upper bound on ζ , or

$$\zeta \leq \frac{(pq\pi - 2c) + (\delta(1 - pq) - \delta\theta)(\tilde{w} - \underline{\tilde{w}}) + (\delta(1 - pq) - \delta\theta)\underline{\tilde{w}} + \delta(1 - pq)\underline{\tilde{v}} - \underline{v}}{c + \delta\theta\tilde{x}(\tilde{w} - \underline{\tilde{w}}) + \delta\theta\underline{\tilde{w}} - \underline{w} - [\delta(1 - pq)\tilde{x}](\tilde{w} - \underline{\tilde{w}})}.$$

We obtain an upper bound on the right side by setting $\tilde{w} - \underline{\tilde{w}} = 0$, obtaining

$$\zeta \leq \frac{(pq\pi - 2c) + (\delta(1 - pq) - \delta\theta)\underline{\tilde{w}} + \delta(1 - pq)\underline{\tilde{v}} - \underline{v}}{c + \delta\theta\underline{\tilde{w}} - \underline{w}}.$$

C.2.2 Front-Loading Effort

We now show that it is impossible for x and \tilde{x} to both be interior. Suppose they are. Then we consider an increase in x and an accompanying decrease in \tilde{x} , effectively moving effort forward. We keep w constant in the process, and show that the result is to increase v , a contradiction.

First, we fix the constraint by differentiating (70) to find

$$\frac{dw}{d\tilde{x}} = \frac{dx}{d\tilde{x}} \frac{w - \underline{w}}{x} + \delta x \theta (\tilde{w} - \underline{\tilde{w}}),$$

and hence, setting $\frac{dw}{d\tilde{x}} = 0$,

$$\frac{dx}{d\tilde{x}} = -\delta x^2 \theta \frac{\tilde{w} - \underline{\tilde{w}}}{w - \underline{w}}. \quad (71)$$

Differentiating (69) and using (71), we have

$$\begin{aligned} \frac{dv}{d\tilde{x}} &= \frac{dx}{d\tilde{x}} \frac{v - \underline{v}}{x} + \delta x ((1 - pq - \theta)[\tilde{w} - \underline{\tilde{w}}] - (1 - pq)[\tilde{v} - \underline{\tilde{v}}]) \\ &= -\delta x \theta \frac{\tilde{w} - \underline{\tilde{w}}}{w - \underline{w}} (v - \underline{v}) + \delta x ((1 - pq - \theta)[\tilde{w} - \underline{\tilde{w}}] - (1 - pq)[\tilde{v} - \underline{\tilde{v}}]). \end{aligned}$$

It concludes the argument to show that this derivative is negative. Multiplying by $w - \underline{w}$, the requisite inequality is

$$(w - \underline{w}) ((1 - pq - \theta)(\tilde{w} - \underline{\tilde{w}}) + (1 - pq)(\tilde{v} - \underline{\tilde{v}})) - (\tilde{w} - \underline{\tilde{w}})(v - \underline{v})\theta < 0.$$

Substituting for $v - \underline{v}$ and $w - \underline{w}$ from (69)–(70) and dropping the common factor x , this is

$$\begin{aligned} & [(1 - pq - \theta)(\tilde{w} - \underline{\tilde{w}}) + (1 - pq)(\tilde{v} - \underline{\tilde{v}})] (c + \delta(\tilde{x}\theta\tilde{w} + (1 - \tilde{x})\theta\underline{\tilde{w}}) - \underline{w}) \\ & < (\tilde{w} - \underline{\tilde{w}})\theta [(1 - \delta)(p\pi - 2c) + \delta(1 - pq - \theta)[\tilde{x}\tilde{w} + (1 - \tilde{x})\underline{\tilde{w}}] \\ & \quad + \delta(1 - pq)(\tilde{x}\tilde{v} + (1 - \tilde{x})\underline{\tilde{v}}) - \underline{v}]. \end{aligned}$$

We can then note that the terms involving \tilde{x} cancel, at which point the expression simplifies to

$$(1 - p - \theta) + (1 - p)\frac{\tilde{v} - \underline{\tilde{v}}}{\tilde{w} - \underline{\tilde{w}}} < \theta \frac{(p\pi - 2c) + (\delta(1 - pq) - \delta\theta)\underline{\tilde{w}} + \delta(1 - pq)\underline{\tilde{v}} - \underline{v}}{c + \delta\theta\underline{\tilde{w}} - \underline{w}},$$

for which, using the definition of ζ , it suffices that $(1 - p - \theta) + (1 - p)\zeta < \theta\zeta$, which is immediate.

An implication of this result is that $x(q)$ is interior for at most one value of q . This in turn implies that the public randomization device is required only for one value of q . If no public randomization device is available, we can construct an equilibrium in which no values of $x(q)$ are interior that approximates the equilibrium examined here. As $\Delta \rightarrow 0$, the public randomization device becomes unimportant and the approximation becomes arbitrarily sharp.

D Online Appendix: Proofs, Observable Effort

D.1 Proof of Proposition 4

We fix $\Delta > 0$, and then suppress the notation for Δ , writing simply δ for $\delta(\Delta) = e^{-r\Delta}$. To capture the effects of delay, we write $\delta\Lambda(q)$ for the effective discounting that elapses before the principal makes an offer at belief q . If the principal undertakes no delay, making the offer as soon as Δ length of time has passed since the previous offer, then $\Lambda(q) = 1$. Delay gives rise to values of $\Lambda(q) < 1$.

We start with the Markov equilibria. As in Bergemann and Hege [2], these raise no issue of existence with observable effort. The usual arguments yield that the agent is either offered no contract, or works on the equilibrium path whenever offered a contract, in which case he is indifferent between doing so or not. We use the same notation as for the unobservable case: v is the principal's payoff, w is the agent's, s is the share, and so on. So the agent's payoff satisfies

$$w(q) = pq\pi(1 - s) + \delta\Lambda(\varphi(q))(1 - pq)w(\varphi(q)) = c + \delta\Lambda(q)w(q), \quad (72)$$

while the principal's payoff solves

$$v(q) = pq\pi s - c + \delta\Lambda(\varphi(q))(1 - pq)v(\varphi(q)). \quad (73)$$

If the project is terminated after one more failure, the values are

$$w(q) = pq\pi(1 - s) = c + \delta\Lambda(q)w(q), \quad v(q) = pq\pi s - c, \quad (74)$$

and so, because the principal is only willing to delay if her payoff is zero, in the last period, combining the equations in (74), either

$$\Lambda(q) = 1, \quad v(q) = pq\pi - \frac{2 - \delta}{1 - \delta}c,$$

or

$$v(q) = 0, \quad \Lambda(q) = \frac{pq\pi - 2c}{\delta(pq\pi - c)}.$$

The first case requires $v(q) \geq 0$ i.e. $q \geq \frac{2-\delta}{1-\delta} \frac{c}{p\pi}$, while the second requires $\Lambda(q) \geq 0$ i.e. $q \geq 2c/(p\pi)$ —a lower threshold. It thus follows that the equilibrium is such that no offer is made for $q \leq \underline{q} := 2c/(p\pi)$, and delay for beliefs q above, but sufficiently close to, \underline{q} .

We shall argue that, at least along any equilibrium path, there is first no delay, and then delay. Let us define as usual the sequence of posterior beliefs, for all $n \geq 0$,

$$q_\tau = \left(1 + \frac{1 - \underline{q}}{\underline{q}}(1 - p)^\tau\right)^{-1}, \quad (75)$$

a sequence of beliefs such that, given q_τ , the effort of an agent takes us to belief $q_{\tau-1}$ (note that $q_0 = \underline{q}$). Let $I_\tau := [q_\tau, q_{\tau+1})$. Fix a Markov equilibrium, and define $\hat{q} := \inf\{q | v(q) > 0\}$ (set $\hat{q} = 1$ if there is no $q \leq 1$ for which $v(q) > 0$) and define $\hat{\tau}$ such that $\hat{q} \in I_{\hat{\tau}}$. We know that $\hat{q} > \underline{q}$. We have, for $\tau = 0$,

$$w(q) = pq\pi - c,$$

and, from (72), for $\tau = 1, \dots, \hat{\tau} - 1$, $q \in I_\tau$, $\tilde{q} = \varphi(q)$,

$$\begin{aligned} w(q) &= pq\pi - c + (1 - pq) \frac{w(\tilde{q}) - c}{\delta w(\tilde{q})} \delta w(\tilde{q}) \\ &= pq(\pi + c) - 2c + (1 - pq)w(\tilde{q}), \end{aligned} \quad (76)$$

where the first equality uses $w(\tilde{q}) = c + \delta\Lambda(\tilde{q})w(\tilde{q})$ to solve for $\Lambda(\tilde{q})$. The solution to this difference equation is

$$w(q) = \pi + c - \frac{2q_0c + (1 - p)^\tau (p(1 - pq_0)\pi + 2c(p(\tau + 1) - (\tau p + 1)q_0))}{p((1 - q_0)(1 - p)^\tau + q_0)}. \quad (77)$$

Taking derivatives, $w'(q)$ is positively proportional to

$$\gamma(\tau) := 2c(1 - p)^{\tau+1} + (1 - (1 - p)^{\tau+1})p\pi + 2c((\tau + 1)p - 1).$$

Because this expression is independent of q , it means, in particular, that the sign of w is constant over each interval I_τ . To evaluate its sign, note that

$$\gamma(\tau + 1) - \gamma(\tau) = 2pc(1 - (1 - p)^{\tau+1}) + p^2(1 - p)^{\tau+1}\pi > 0,$$

so that, if w is increasing on I_τ , it is also increasing on $I_{\tau+1}$. Because it is increasing on I_0 , it is increasing on each interval.

Consider now some \check{q} arbitrarily close to \hat{q} such that $v(\check{q}) > 0$. Then $\Lambda(\check{q}) = 1$, and so $w(\check{q}) = c/(1 - \delta)$. Note that, because $v(\hat{q}) = 0$, we can write, for all beliefs $q \in I_{\hat{\tau}} \cap [\check{q}, 1]$, using (73) first, and then (72),

$$v(q) = pq\pi - c - pq(1 - s)\pi = pq\pi - c + \delta(1 - pq)\Lambda(\check{q})w(\check{q}) - w(q).$$

The term $pq\pi - c + \delta(1 - pq)\Lambda(\check{q})w(\check{q})$ must be increasing in q : it is precisely the definition of $w(q)$ in the sequence studied above.³⁶ The last term, $-w(q)$, is minimized at \check{q} , since it equals $-c/(1 - \delta)$ there. Therefore, v must be also strictly positive for all beliefs $q \in I_{\hat{\tau}} \cap [\check{q}, 1]$, and both v, w must be continuous at \hat{q} . This means that $(\hat{\tau}, \hat{q})$ are such that, for some $q_0 \in I_0$, (77) holds with $w(\hat{q}) = c/(1 - \delta)$. Note that (76) gives that, for $q = \hat{q}$,

$$\frac{c}{1 - \delta} = p\hat{q}\pi - 2c + p\hat{q}c + (1 - p\hat{q})w(\check{q}),$$

and so, since $w(\check{q}) < c/(1 - \delta)$, $\delta p\hat{q}c < (1 - \delta)(p\hat{q}\pi - 2c)$, or equivalently

$$((1 - \delta)p\pi - \delta pc)\hat{q} > 2(1 - \delta)c,$$

which implies that, at the very least, $p\pi - \delta pc/(1 - \delta) > 0$ (from which it is apparent that the existence of such a $\hat{q} < 1$ only holds for some parameters). Consider now the belief q such that $\tilde{q} = \hat{q}$. If $\Lambda(q) = 1$, then, solving for $s(q)$ by using $w(q) = w(\tilde{q}) = c/(1 - \delta)$, we get

$$\frac{v(q)}{1 - q} = pq\pi - 2c - \frac{\delta pqc}{1 - \delta} + \delta \frac{v(\tilde{q})}{1 - \tilde{q}}.$$

Because, as we have seen, $p\pi - \delta pc/(1 - \delta) > 0$, the term $pq\pi - 2c - \frac{\delta pqc}{1 - \delta}$ is increasing in q , and since it is non-negative at \hat{q} , it is strictly positive at q . Therefore, $v(q) > 0$, and it is clear that there cannot be delay at q , because $\Lambda(q) < 1$ would imply a higher value of $s(q)$, and thus $v(q)$ would still be strictly positive. Indeed, this argument applies to any q for which $\tilde{q} \geq \hat{q}$ and $w(\tilde{q}) = c/(1 - \delta)$. This implies that, for any sequence of beliefs that can be obtained from Bayes' rule after strings of failures, the equilibrium must be such that v is first strictly positive (when the belief is high enough, and the prior might not be enough to begin with), after which $v = 0$ and there is delay until the belief drops below \underline{q} at which point the project is abandoned.

³⁶Of course, this is not the value of the agent at q , since now $\Lambda = 1$.

This does not, however, imply that $v(q) = 0$ if and only if $q < \hat{q}$. The discreteness of the problem does not rule out multiple solutions to (77). It remains to show that all such solutions converge to the same belief as period length shrinks. Replace p , c , $1 - \delta$ by $p\Delta$, $c\Delta$ and $r\Delta$ respectively, and let $\kappa = \tau\Delta$. Taking limits in (77), we obtain that the value of κ for which $\Lambda(q) = 1$, i.e. $w(q) = c/(1 - \delta)$, solves

$$e^{\kappa r \sigma} = \left(1 + \frac{\sigma}{2} + \kappa r \sigma\right) \frac{\psi}{\psi - \sigma}, \quad (78)$$

so all solutions \hat{q} converge to the same solution q^* as $\Delta \rightarrow 0$. Taking the same limits in (75), the corresponding belief threshold q^* solves

$$e^{\kappa r \sigma} = \frac{\psi}{2} \frac{q^*}{1 - q^*}.$$

Substituting into (78), and solving, gives that

$$q^* = 1 - \frac{1}{1 - 2 \frac{W_{-1}\left(-\frac{\psi - \sigma}{\psi} e^{-1 - \frac{\sigma}{2}}\right)}{\psi - \sigma}},$$

where W_{-1} is the negative branch of the Lambert function (the positive branch only admits a solution to the equation that is below \underline{q}). Then $q^* < 1$ if and only if $\xi > 1 + \sigma$. Otherwise, as $\Delta \rightarrow 0$, $v(q) = 0$ for all $q \in [\underline{q}, 1]$.

D.2 Proof of Proposition 5

Let us assume throughout that

$$1 - \delta \leq \frac{1}{\frac{p\pi}{c} - 1},$$

which is automatically satisfied as $\Delta \rightarrow 0$, since the left side is approximately $r\Delta$, while the right side converges to the positive constant $1/(1 + \psi)$.

We start by arguing that equilibria in which the principal makes zero profits exist for every $q < 1$. If such an equilibrium exists, then there is a “full-stop” equilibrium in which the project is terminated at this belief, i.e. the principal offers no contract, with the threat that doing otherwise would lead to reversion to the equilibrium in which the principal makes zero profits. Let \tilde{q} denote the infimum over values of q for which such an equilibrium does not exist. From the analysis of Markov equilibria, we know that $\tilde{q} > \underline{q}$. Consider some q above \tilde{q} for which it does not exist, and such that a failure leads to a belief strictly below \tilde{q} . That is, we can specify that the game terminates after a failure. To

see whether there exists an equilibrium in which the principal makes zero profits starting at q , we solve (74), which gives as necessary and sufficient condition that

$$\Lambda(q) = \frac{pq\pi - 2c}{\delta(pq\pi - c)} \in [0, 1],$$

which follows from our assumption on δ . This is the desired contradiction: a full-stop equilibrium exists for all values of q .

The best equilibrium for the principal, then, obtains if cheating by the agent is threatened by termination. Setting $\Lambda(q) = 1$ is then optimal, unless it is best to terminate the project. The agent prefers to work at the last stage (and thus, at all stages) if and only if

$$pq\pi(1 - s) \geq c,$$

so that the seller's payoff at the last stage is

$$v(q) = pq\pi - 2c, \text{ and } s(q) = 1 - \frac{c}{pq\pi},$$

and so the project is terminated as soon as the posterior belief drops below $\underline{q} = 2c/(p\pi)$. More generally, the values are obtained from solving

$$w(q) = pq(1 - s)\pi + \delta(1 - pq)w(\varphi(q)) = c, v(q) = pqs\pi - c + \delta(1 - pq)v(\varphi(q)),$$

from which we get

$$\frac{v(q)}{1 - q} = \frac{pq(\pi - \delta c) - (2 - \delta)c}{1 - q} + \delta \frac{v(\varphi(q))}{1 - \varphi(q)}, \quad (79)$$

except when $q \in I_0$, when $v(q) = pq\pi - 2c$, and so it is optimal to terminate as soon as the belief drops below \underline{q} . Equation (79) is straightforward to solve explicitly, and taking limits gives the value given by (30).

E Online Appendix: Derivations and Proofs, Good Projects

This appendix examines the case in which the project is known to be good ($\bar{q} = 1$). We fix $\Delta > 0$, but omit Δ from the notation whenever we can do so without confusion.

E.1 The First-Best Policy

The value of conducting an experiment is given by

$$\begin{aligned} V &= p\pi - c + \delta(1 - p)V \\ &= \frac{(p\pi - c)}{1 - \delta(1 - p)}. \end{aligned}$$

The optimal action is to experiment if and only if $V \geq 0$, or

$$p \geq \frac{c}{\pi}. \quad (80)$$

The first-best strategy thus either never conducts any experiments, or relentlessly conducts experiments until a success is realized, depending on whether $p < c/\pi$ or $p > c/\pi$.

E.2 Stationary No-Delay Equilibrium: Impatient Projects

We first investigate Markov equilibria. We begin with a candidate equilibrium in which the principal extends funding at every opportunity, and the agent exerts effort in each case. If the principal offers share s , she receives an expected payoff in each period of

$$p\pi s - c.$$

The agent's payoff solves, by the principle of optimality,

$$\begin{aligned} W &= \max\{c + \delta W, p\pi(1-s) + \delta(1-p)W\} \\ &= \max\left\{\frac{c}{1-\delta}, \frac{p\pi(1-s)}{1-\delta(1-p)}\right\}. \end{aligned} \quad (81)$$

Such an equilibrium will exist if and only if the principal finds it optimal to fund the project and the agent finds it optimal to work, or

$$\begin{aligned} p\pi s &\geq c, \\ \frac{p\pi(1-s)}{1-\delta(1-p)} &\geq \frac{c}{1-\delta}. \end{aligned}$$

Combining and rearranging, this is equivalent to

$$p \cdot \min\{(1-\delta)\pi s, (1-\delta)\pi(1-s) - \delta c\} \geq (1-\delta)c.$$

There is some value of $s \in [0, 1]$ rendering the second term in the minimum positive, a necessary condition for the agent to work, only if $(1-\delta)\pi > \delta c$. If this is the case, then since the arguments of the minimum vary in opposite directions with respect to s , the lowest value of p or lowest ratio π/c for which such an equilibrium exists is attained when the two terms are equal, that is, when

$$s = \frac{1}{2} \left(1 - \delta \frac{c}{(1-\delta)\pi} \right), \quad (82)$$

in which case the constraint reduces to

$$\frac{\pi}{c} \geq \frac{2}{p} + \frac{\delta}{1-\delta}, \quad (83)$$

which implies $(1 - \delta)\pi > \delta c$. Hence, necessary and sufficient conditions for the existence of a full-effort stationary equilibrium are that the players be sufficiently impatient to satisfy (83). Taking the limit as $\Delta \rightarrow 0$, the constraint given by (83) becomes

$$\psi > \sigma, \tag{84}$$

which we have deemed impatient projects.

The principal will choose s to make the agent indifferent between working and shirking, giving equality of the two terms in (81) and hence an agent payoff of $W^* = c/(1 - \delta)$. This is expected—by always shirking, the agent can secure a payoff of c . In a Markov equilibrium, this must also be his unique equilibrium payoff, since the principal has no incentive to offer him more than the minimal share that induces him to work (the continuation play being independent of current behavior).

The total surplus S of the project satisfies

$$S = p\pi - c + \delta(1 - p)S, \quad \text{or} \quad S = \frac{(p\pi - c)}{1 - \delta(1 - p)}.$$

The principal's payoff is then

$$\frac{(p\pi - c)}{1 - \delta(1 - p)} - \frac{c}{1 - \delta} = \frac{(1 - \delta)(p\pi - 2c) - \delta pc}{(1 - \delta)(1 - \delta(1 - p))} =: V^*,$$

which, in the limit as $\Delta \rightarrow 0$, is positive if and only if $\psi > \sigma$.

E.3 Markov Equilibria for Other Parameters

E.3.1 Patient Projects: Delay

It is straightforward that there is no equilibrium with experimentation if $p\pi - 2c < 0$. We accordingly consider the remaining case in which

$$\frac{2}{p} < \frac{\pi}{c} < \frac{2}{p} + \frac{\delta}{1 - \delta},$$

or, in the limit as $\Delta \rightarrow 0$,

$$0 < \psi < \sigma,$$

giving a patient project.

We now have an equilibrium with delay. The principal waits $\Delta\Psi$ time between offers, with $\Psi \geq 1$. The agent exerts effort at each opportunity, but is indifferent between doing so and shirking, and so his payoff is $\Delta(c + \delta(\Delta\Psi)c + \delta(\Delta\Psi)^2c + \dots) = \frac{c\Delta}{1 - \delta(\Delta\Psi)}$.³⁷

³⁷More formally, the agent's strategy specifies that he works if and only if $s \geq c/(p\pi)$.

The principal is indifferent in each period between offering the contract $s < 1$ and delaying such an offer, and so it must be that she just breaks even: $ps\pi = c$. On the other hand, since the agent is indifferent between shirking and not, we must have

$$c\Delta + \delta(\Delta\Psi)\frac{c\Delta}{1 - \delta(\Delta\Psi)} = p\Delta(1 - s)\pi + \delta(\Delta\Psi)(1 - p\Delta)\frac{c\Delta}{1 - \delta(\Delta\Psi)}.$$

Using $s = c/p\pi$, this gives

$$\frac{\Delta\delta(\Delta\Psi)}{1 - \delta(\Delta\Psi)} = \frac{\pi}{c} - \frac{2}{p}.$$

Using the approximation $\delta(\Delta\Psi) = 1 - r\Delta\Psi$ (for small Δ), in the limit as $\Delta \rightarrow 0$, we have

$$\frac{1}{\Psi} = r \left(\frac{\pi}{c} - \frac{2}{p} \right).$$

Delay is thus zero (i.e., $\Psi = 1$) when $\frac{\pi}{c} = \frac{2}{p} + \frac{1}{r}$, and increases without bound as ψ approaches zero.

The payoff of the principal in this equilibrium is 0, and the agent's payoff is

$$W = \frac{p\pi - 2c}{\delta p}.$$

We now have completed the characterization of Markov equilibria, yielding payoffs that are summarized in Figure 5.

E.4 Non-Markov Equilibria

We now extend our analysis to a characterization of all equilibria. We first find equilibria with stationary outcomes backed up by the threat of out-of-equilibrium punishments, and then use these to construct a family of equilibria with nonstationary outcomes.

Our first step is the following lemma, proved in Appendix E.5.1.

Lemma 19 *The agent's equilibrium payoff never exceeds $\frac{c}{1-\delta}$.*

E.4.1 Impatient Projects

Suppose first that $\frac{\pi}{c} \geq \frac{2}{p} + \frac{\delta}{1-\delta}$. Section E.2 established that there then exists a Markov equilibrium in which the agent always works on the equilibrium path, with payoffs

$$(W^*, V^*) := \left(\frac{c}{1-\delta}, \frac{(1-\delta)(p\pi - 2c) - \delta pc}{(1-\delta)(1-\delta(1-p))} \right).$$

It is immediate that V^* puts a lower bound on the principal's payoff in *any* equilibrium. In particular, the share s offered by the principal in this equilibrium necessarily induces




Figure 5: Payoffs from the Markov equilibrium of a project known to be good ($\bar{q} = 1$), as a function of the “benefit-cost” ratio π/c , fixing c (so that we can identify c on the vertical axis). Both players obviously earn zero in the null equilibrium of an unprofitable project. The principal’s payoff is fixed at zero for patient projects, while the agent’s increases as does π . The agent’s payoff is fixed at $c\Delta/(1 - \delta(\Delta))$ for impatient projects, while the principal’s payoff increases in π .

the agent to work, since it does so when the agent expects his maximum continuation payoff of W^* (cf. Lemma 19), and hence when it is hardest to motivate the agent. By continually offering this share, the principal can then be assured of payoff V^* .

We begin our search for additional equilibrium payoffs by constructing a family of potential equilibria with stationary equilibrium paths. We assume that after making an offer, the principal waits a length of time $\Delta\Psi$ until making the next offer, where $\Psi \geq 1$.

Why doesn’t the principal make an offer to the agent as soon as possible? Doing so prompts an immediate switch to the full-effort equilibrium with payoffs (W^*, V^*) (with the agent shirking unless offered a share at least as large as in the full-effort equilibrium). We will then have an equilibrium as long as the principal’s payoff exceeds $V^*/\delta((\Psi - 1)\Delta)$, ensuring that the principal would rather wait an additional length of time $(\Psi - 1)\Delta$ to continue along the equilibrium path, rather than switch immediately to the full-effort equilibrium.

The agent is indifferent between working and shirking, whenever offered a nontrivial contract, and so his payoff is $c\Delta + \delta(\Delta\Psi)c\Delta + \dots = \frac{c\Delta}{1 - \delta(\Delta\Psi)}$. Using this continuation value, the agent’s incentive constraint is

$$p(1 - s)\pi\Delta + \delta(\Delta\Psi)(1 - p\Delta)\frac{c\Delta}{1 - \delta(\Delta\Psi)} = c\Delta + \delta(\Delta\Psi)\frac{c\Delta}{1 - \delta(\Delta\Psi)},$$

or

$$(p\pi - 2c)\Delta - \delta(\Delta\Psi)p\Delta\frac{c\Delta}{1 - \delta(\Delta\Psi)} = (p\pi s - c)\Delta.$$

Using this for the second equality, the principal's value is then

$$\begin{aligned} V &= (ps\pi - c)\Delta + \delta(\Delta\Psi)(1 - p)V \\ &= (p\pi - 2c)\Delta - \delta(\Delta\Psi)p\Delta\frac{c\Delta}{1 - \delta(\Delta\Psi)} + \delta(\Delta\Psi)(1 - p\Delta)V \\ &= \frac{(1 - \delta(\Delta\Psi))(p\pi - 2c)\Delta - \delta(\Delta\Psi)pc\Delta^2}{(1 - \delta(\Delta\Psi))(1 - \delta(\Delta\Psi)(1 - p\Delta))}. \end{aligned}$$

This gives us a value for the principal that equals V^* when $\Psi = 1$, in which case we have simply duplicated the stationary full-effort equilibrium. However, these strategies may give equilibria with a higher payoff to the principal, and a lower payoff to the agent, when $\Psi > 1$. In particular, as we increase Ψ , we decrease both the total surplus and the rent that the agent can guarantee by shirking. This implies that the principal might be better off slowing down the project from $\Psi = 1$, if the cost of the rent is large relative to the profitability of the project, i.e., if π/c is relatively low. Indeed, this returns us to the intuition behind the existence of Markov equilibria with delay for low-discount projects, where π/c is too low for the existence of an equilibrium with $\Psi = 1$: by slowing down the project, the cost of providing incentives to the agent is decreased, and hence the principal's payoff might increase.³⁸

Let $V(\Psi)$ denote the principal's payoff as a function of Ψ . We have $\lim_{\Psi \rightarrow \infty} V(\Psi) = 0$, giving the expected result that there is no payoff when no effort is invested. Are there any values for which $V(\Psi) > V(1)$? The function $V(\cdot)$ is concave, and the function $V(\Psi) = V(1)$ admits a unique root $\Psi^\dagger > 1$ if its derivative at 1 is positive. It is most convenient to examine the limiting case in which $\Delta \rightarrow 0$, allowing us to write

$$V = \frac{r\Psi(p\pi - 2c) - pc}{r\Psi(r\Psi + p)}$$

and then to note that the resulting derivative in Ψ has numerator equal to

$$r\Psi(r\Psi + p)(p\pi - 2c) - [r\Psi(p\pi - 2c) - pc](2r^2\Psi + rp).$$

Taking $\Psi = 1$, this is positive if

$$r(r + p)(p\pi - 2c) > [r(p\pi - 2c)](2r + p),$$

³⁸We were considering Markov equilibria when examining patient projects, and hence the optimality of delay required that the principal be indifferent between offering a contract and not offering one, which in turn implied that the principal's payoff was fixed at zero. Here, we are using the nonstationary threat of a punishment to payoff V^* to enforce the delay, and hence the principal need not be indifferent and can earn a positive payoff.

which simplifies to

$$\frac{\pi}{c} < \frac{2}{p} + \frac{1}{r} \left(2 + \frac{p}{r}\right).$$

We must then split our analysis of impatient projects into two cases. If π/c is large (i.e., $\frac{\pi}{c} > \frac{2}{p} + \frac{1}{r} \left(2 + \frac{p}{r}\right)$), then $V(\Psi) < V^*$ for all $\Psi > 1$. Therefore, our search for non-Markov equilibria has not yet turned up any additional equilibria. Indeed, Lemma (20) shows that there are *no* other equilibria in this case. Alternatively, if π/c is not too large ($\frac{2}{p} + \frac{1}{r} \leq \frac{\pi}{c} < \frac{2}{p} + \frac{1}{r} \left(2 + \frac{p}{r}\right)$), then as the delay factor Ψ rises above unity, the principal's payoff initially increases. We have then potentially constructed an entire family of stationary-outcome equilibria, one for each value $\Psi \in [1, \Psi^\dagger]$ (recalling again that V is concave).³⁹ These nonstationary (but stationary-outcome) equilibria give the agent a payoff less than $W^* = c/(1 - \delta)$ and the principal a payoff larger than V^* .

The following lemma, proven in Appendix E.5.2, states that these equilibria yield the lowest equilibrium payoff to the agent.

Lemma 20

[20.1] [Very Impatient Projects] If

$$\frac{\pi}{c} \geq \frac{2}{p} + \frac{\delta}{1 - \delta} \left(2 + \frac{\delta}{1 - \delta} p\right),$$

then the lowest equilibrium payoff \underline{W} to the agent is given by $W^* = \frac{c}{1 - \delta}$. Hence, there is then a unique equilibrium with payoffs (W^*, V^*) .

[20.2] [Moderately Impatient Projects] If

$$\frac{2}{p} + \frac{\delta}{1 - \delta} \leq \frac{\pi}{c} < \frac{2}{p} + \frac{\delta}{1 - \delta} \left(2 + \frac{\delta}{1 - \delta} p\right),$$

then the infimum over equilibrium payoffs to the agent (as $\Delta \rightarrow 0$) is given by $W(\Psi^\dagger) = \frac{V^*}{\delta p} \leq \frac{c}{1 - \delta}$.

In the latter case, the limit of the equilibria giving the agent his lowest equilibrium payoff, as $\Delta \rightarrow 0$, sets $\Psi = \Psi^\dagger$ and gives the principal payoff V^* , and so gives both players their lowest equilibrium payoff. To summarize these relationships, it is convenient let

$$(W(\Psi^\dagger), V^*) = (W(\Psi^\dagger), V(\Psi^\dagger)) := (\underline{W}, \underline{V}).$$

We have now established (W^*, V^*) as the unique equilibrium payoffs for very impatient projects. For moderately impatient projects, we have bounded the principal's payoff below by \underline{V} and bounded the agent's payoff below by \underline{W} and above by W^* .

³⁹We have an equilibrium for each $\Psi \in [1, \Psi^\dagger]$ satisfying the incentive constraint $\delta(\Psi - 1)\Delta V(\Psi) \geq V^*$. As $\Delta \rightarrow 0$, the set of such z converges to the entire interval $[1, \Psi^\dagger]$.

To characterize the complete set of equilibrium payoffs for moderately impatient projects, we must consider equilibria with nonstationary outcomes. Appendix E.5.3 establishes the following technical lemma:

Lemma 21 *Let the parameters satisfy $\frac{2}{p} + \frac{\delta}{1-\delta} \leq \frac{\pi}{c} < \frac{2}{p} + \frac{\delta}{1-\delta} (2 + \frac{\delta}{1-\delta}p)$ and let (W, V) be an arbitrary equilibrium payoff. Then*

$$\frac{V - \underline{V}}{W - \underline{W}} \leq \frac{\delta p}{1 - \delta} = \frac{\underline{V}}{\underline{W}}.$$

The geometric interpretation of this lemma is immediate: the ratio of the principal's to the agent's payoff is maximized by the limiting worst payoffs $(\underline{W}, \underline{V})$.

Any equilibrium payoff can be achieved by an equilibrium in which, in the first period, the equilibrium delivering the worst equilibrium payoff to the agent is played with some probability $1 - x_0$, and an extremal equilibrium (i.e, an equilibrium with payoffs on the boundary of the set of equilibrium payoffs) is played with probability x_0 .⁴⁰ If the former equilibrium is chosen, subsequent play continues with that equilibrium. If the latter equilibrium is chosen, then the next period again features a randomization attaching probability $1 - x_1$ to an equilibrium featuring the worst possible payoff to the agent and attaching probability x_1 to an extremal equilibrium. Continuing in this way, we can characterize an equilibrium giving payoffs (W_0, V_0) as a sequence $\{x_t, (W_t, V_t)\}_{t=0}^\infty$, where x_t is the probability that an equilibrium with the extremal payoffs (W_t, V_t) is chosen in period t , conditional on no previous mixture having chosen the equilibrium with the worst equilibrium payoffs to the agent.

Given W , consider the supremum over values of V among equilibrium payoffs, and say that the resulting payoff (W, V) is on the frontier of the equilibrium payoff set. Our goal is to characterize this frontier. If (W_0, V_0) is on the frontier, it sacrifices no generality to assume that in each of the equilibria yielding payoffs (W_t, V_t) (in the sequence $\{x_t, (W_t, V_t)\}_{t=0}^\infty$), the principal offers a contract to the agent without delay.⁴¹ In addition, we can assume that each such equilibrium calls for the principal to offer some share s_t

⁴⁰Because the set of equilibrium payoffs is bounded and convex, any equilibrium payoff can be written as a convex combination of two extreme payoffs. One of these extreme payoffs can be chosen freely, and hence can be taken to feature the worst equilibrium payoff.

⁴¹If the principal delays, we can view the resulting equilibrium payoff as a convex combination of two payoffs, one (denoted by (W'_t, V'_t)) corresponding to the case in which a contract is offered and one corresponding to offering no contract. But the latter is an interior payoff of the form $(\delta(W''_t, V''_t))$, given by δ times the accompanying continuation payoff (W''_t, V''_t) . We can then replace (W_t, V_t) by a convex combination of (W'_t, V'_t) and (W''_t, V''_t) , to obtain a payoff of the form (W_t, V_t^\dagger) , with $V_t^\dagger > V_t$. Because W_t is unchanged, none of the incentives in previous periods are altered ensuring that we still have an equilibrium. Because the principal's payoff V_t has increased, so has V_0 , contradicting the supposition that the latter was on the payoff frontier.

to the agent that induces the agent to work.⁴² Using Lemma 21, Appendix E.5.4 proves the following, completing our characterization of the equilibrium frontier in the case of an impatient project:

Lemma 22 *In an equilibrium whose payoff is on the frontier of the equilibrium payoff set, it cannot be that both $x_t \in (0, 1)$ and $x_{t+1} \in (0, 1)$. More precisely, x_t is weakly decreasing in t , and there is at most one value of t for which x_t is in $(0, 1)$.*

This lemma tells us that the equilibria on the frontier can be described as follows: for some $T \in \mathbb{N} \cup \{\infty\}$ periods, the project is funded without delay by the principal, and the agent exerts effort, being indifferent between doing so or not. From period T onward, an equilibrium giving the agent his worst payoff is played. We have already seen the two extreme points of this family: if $T = \infty$, there is never any delay, resulting in the payoff pair (W^*, V^*) . If $T = 0$, the worst equilibrium is obtained. For very impatient projects, all these equilibria are equivalent (since the no-delay equilibrium *is* then the worst equilibrium), and only the payoff vector (W^*, V^*) is obtained. For moderately impatient projects, however, this defines a sequence of points (one for each possible value of T), the convex hull of which defines the set of all equilibrium payoffs. Any payoff in this set can be achieved by an equilibrium that randomizes in the initial period between the worst equilibrium, and an equilibrium on the frontier.

This result in turn leads to a concise characterization of the set of equilibrium payoffs, in the limit as $\Delta \rightarrow 0$. In particular, as $\Delta \rightarrow 0$, the set of equilibrium payoffs converges to a set bounded below by the line segment connecting the payoffs $(\underline{W}, \underline{V})$ and $(W^*, \underline{V}) = (W^*, V^*)$, and bounded above by a payoff frontier characterized by Lemma 22. This set of payoffs is illustrated in the two right panels of Figure 6. An analytical determination of the set of equilibrium payoffs is provided in Section E.4.3, for the convenient case in which the length of a time period Δ is arbitrarily small.

E.4.2 Patient Projects

Consider now the case in which

$$\frac{2}{p} \leq \frac{\pi}{c} < \frac{2}{p} + \frac{\delta}{1 - \delta}.$$

The Markov equilibria in this region involve a zero payoff for the principal. This means, in particular, that we can construct an equilibrium in which both players' payoff is zero: on the equilibrium path, the principal makes no offer to the agent; if she ever deviates, both

⁴²Should the principal be called upon to offer a contract that induces the agent to shirk, it is a straightforward calculation that it increases the principal's payoff, while holding that of the agent constant, to increase the share s_t just enough to make the agent indifferent between working and shirking, and to have the agent work, again ensuring that (W, V) is not extreme.

players play the stationary equilibrium from that point on, which for those parameters also yields zero profit to the principal. Since this equilibrium gives both players a payoff of zero, it is trivially the worst equilibrium.

Lemma 22 is valid here as well,⁴³ and so the equilibrium payoffs on the frontier are again obtained by considering the strategy profiles indexed by some integer T such that the project is funded for the first T periods, and effort is exerted (the agent being indifferent doing so), after which the worst equilibrium is played. Unlike in the case of an impatient project, we now have a constraint on T . In particular, as $T \rightarrow \infty$, the value to the principal of this strategy profile becomes negative. Since the value must remain nonnegative in equilibrium, this defines an upper bound on the values of T that are consistent with equilibrium. While the sequence of such payoffs can be easily computed, and the upper bound implicitly defined, the analysis is once again crisper when we consider the continuous-time limit $\Delta \rightarrow 0$, as in Section E.4.3. The set of equilibrium payoffs is illustrated on the second panel of Figure 6.

E.4.3 Characterization of Equilibrium Payoffs

Sections E.4.1–E.4.2 characterize the set of equilibrium payoffs. However, this characterization is not easy to use, as the difference equations describing the boundaries of the equilibrium payoff set are rather unwieldy. We consider here the limit of these difference equations, and hence of the payoff set, as we let the length Δ of a period tend to 0.

Given an equilibrium in which there is no delay the agent invariably exerts effort, the value V_t at time t to the principal solves (up to terms of order Δ^2 or higher)

$$V_t = p\pi s_t \Delta - c\Delta + (1 - (r + p)\Delta)(V_t + \dot{V}_t \Delta),$$

or, in the limit as $\Delta \rightarrow 0$,

$$0 = p\pi s_t - c - (r + p)v(t) + \dot{v}(t), \quad (85)$$

where s_t is the share to the principal in case of success, and \dot{v} is the time derivative of v (whose differentiability is easy to derive from the difference equations). Similarly, if the agent is indifferent between exerting effort or not, we must have (up to terms of order Δ^2 or higher)

$$W_t = p\pi(1 - s_t)\Delta + (1 - (r + p)\Delta)(W_t + \dot{W}_t \Delta) = c\Delta + (1 - r\Delta)(W_t + \dot{W}_t \Delta),$$

where W_t is the agent's continuation payoff from time t onwards. In the limit as $\Delta \rightarrow 0$, this gives

$$0 = p\pi(1 - s_t) - (r + p)w_t + \dot{w}_t = c - rw(t) + \dot{w}_t. \quad (86)$$

⁴³In this range of parameters, $\underline{W} = \underline{V} = 0$, and upon inserting these values, the proof of Lemma 22 continues to hold. From (90), the counterpart of Lemma 21 in this case is $\frac{\underline{V}}{\underline{W}} \leq \frac{p\pi - 2c}{c}$.

We may use these formulae to obtain closed-forms in the limit for the boundaries of the payoff sets described above.

Let us first ignore the terminal condition and study the stationary case in which $\dot{v}_t = \dot{w}_t = 0$ for all t . Then

$$w_t = w^* := \frac{c}{r}, \quad v_t = v^* := \frac{\psi - \sigma c}{\sigma + 1} \frac{c}{r},$$

which are positive provided $\psi \geq \sigma$. If instead $\psi < \sigma$, the principal's payoff is zero in the unique stationary equilibrium. It is easy to check that if in addition $\psi < 0$, it is not possible to have the agent exert effort in any equilibrium, and the unique equilibrium payoff vector is $(0, 0)$. This provides us with two of the relevant boundaries, between unprofitable and patient projects, and between patient and moderately impatient projects. The derivation of the boundary between moderately impatient and very impatient projects is more involved, and available along with the proof of Proposition 9 in Section E.5.5.

Proposition 9 *The set of equilibrium payoffs for a project that is known to be good ($q = 1$), in the limit as period length becomes short, is given by:*

- **Unprofitable Projects** ($\psi < 0$). No effort can be induced, and the unique equilibrium payoff is $(w, v) = (0, 0)$.
- **Patient Projects** ($0 < \psi < \sigma$). The set of equilibrium payoffs is given by the pairs (w, v) , where $w \in [0, w^\dagger]$, and

$$0 \leq v \leq \frac{\psi + 1}{\sigma + 1} \left[1 - \left(1 - \frac{w}{c} \right)^{\sigma+1} \right] \frac{c}{r} - w,$$

where w^\dagger is the unique positive value for which the upper extremity of this interval is equal to zero. In the equilibria achieving payoffs on the frontier, there is no delay, and the agent always exerts effort, until some time $T < \infty$ at which funding stops altogether. Such equilibria exist for all T below some parameter-dependent threshold \bar{T} .

- **Moderately Impatient Projects** ($\sigma < \psi < \sigma(\sigma + 2)$). The set of equilibrium payoffs is given by the pairs (w, v) , for $w \in [\underline{w}, \frac{c}{r}]$, and

$$v^* \leq v \leq \frac{c}{r} \left(\frac{\psi + 1}{\sigma + 1} + \frac{\psi - \psi\sigma - 2\sigma}{\psi - \sigma^2 - 2\sigma} \left(\frac{\psi - \sigma^2}{\sigma^2 + \sigma} \right)^{-\sigma} \left(1 - \frac{rw}{c} \right)^{\sigma+1} \frac{c}{r} \right) - w,$$

where $v^* = \frac{\psi - \sigma c}{\sigma + 1} \frac{c}{r}$ and $\underline{w} = v^*/\sigma$. In the equilibria achieving payoffs on the frontier, there is no delay, and the agent exerts effort, until some time $T \leq \infty$ from which point on there is delay, with continuation payoff (\underline{w}, v^*) .

- **Very Impatient Projects** ($\psi > \sigma(\sigma + 2)$). The unique equilibrium payoff involves no delay and the agent exerting effort: $(w, v) = (w^*, v^*) = \left(\frac{c}{r}, \frac{\psi - \sigma c}{\sigma + 1} \frac{c}{r} \right)$.




Figure 6: Set of equilibrium payoffs for a project that is known to be good ($\bar{q} = 1$), for the limiting case of arbitrarily short time periods ($\Delta \rightarrow 0$). We measure the agent's payoff w on the horizontal axis and the principal's payoff v on the vertical axis. To obtain concrete results, we set $c/r = p/r = 1$ and, from left to right, $(p\pi - c)/c = 0$ (unprofitable project), $(p\pi - c)/c = 3/2$ (patient project), $(p\pi - c)/c = 3$ (moderately impatient project), and $(p\pi - c)/c = 7$ (very impatient project). The point in each case identifies the payoffs of Markov equilibria. The dotted line in the case of a moderately impatient project identifies the payoffs of the equilibria with stationary outcomes, and the shaded areas identify the sets of equilibrium payoffs. Note that neither axis in the third panel starts at 0.

E.4.4 Summary

Figure 6 summarizes our characterization of the set of equilibrium payoffs, for the limiting case as $\Delta \rightarrow 0$. In each case, the Markov equilibrium puts a lower bound on the principal's payoff. For either very impatient or (of course) unprofitable projects, there are no other equilibria. It is not particularly surprising that, for moderately impatient projects, there are equilibria with stationary outcomes backed up by out-of-equilibrium punishments that increase the principal's payoff. The principal has a commitment problem, preferring to reduce the costs of current incentives by reducing the pace and hence the value of continued experimentation. The punishments supporting the equilibrium path in the case of moderately impatient projects effectively provide such commitment power, allowing the principal to increase her payoff at the expense of the agent. It is somewhat more surprising that for patient and moderately impatient projects the principal's payoff is maximized by an equilibrium whose outcome is nonstationary, coupling an initial period of no delay with a future in which there is either delay or the project is altogether. Moreover, in the case of a patient project, such equilibria can increase the payoffs of both agents.

E.5 Proofs

E.5.1 Proof of Lemma 19

Let \bar{W} be the agent's maximal equilibrium payoff. We can restrict attention to cases in which the principal has offered a contract to the agent, and in which the agent works.⁴⁴

We first note that a lower bound on the principal's payoff is provided by *always* choosing that value s^W satisfying (and hence inducing the agent to work, no matter how lucrative a continuation value the agent expects)

$$p\pi(1 - s^W) + \delta(1 - p)\bar{W} = c + \delta\bar{W},$$

which we can rearrange to give

$$p\pi s^W - c = -\delta p\bar{W} + p\pi - 2c,$$

and hence a principal payoff of

$$\frac{ps^W\pi - c}{1 - \delta(1 - p)} = \frac{p\pi - 2c - \delta p\bar{W}}{1 - \delta(1 - p)}.$$

We can then characterize \bar{W} as the solution to the maximization problem:

$$\begin{aligned} \bar{W} &= \max_{s,W,V} p\pi(1 - s) + \delta(1 - p)W \\ \text{s.t. } &\bar{W} \geq c + \delta W, \\ &\bar{W} \geq W, \\ &p\pi s - c + \delta(1 - p)V \geq \frac{p\pi - 2c - \delta p\bar{W}}{1 - \delta(1 - p)}, \\ &V + W \leq \frac{p\pi - c}{1 - \delta(1 - p)}, \end{aligned}$$

where the first constraint is the agent's incentive constraint, the second establishes \bar{W} as the largest agent payoff, the third imposes the lower bound on the principal's payoff, and the final constraint imposes feasibility. Notice that if the first constraint binds, then (using the second constraint) we immediately have $\bar{W} \leq \frac{c}{1-\delta}$, and so we may drop the first constraint. Next, the final constraint will surely bind (otherwise we can decrease s and increase V so as to preserve the penultimate constraint while increasing the objective),

⁴⁴If $c/(1 - \delta)$ is an upper bound on the agent's payoff conditional on a contract being offered, then it must also be an upper bound on an equilibrium path in which a contract is offered only after some delay. Next, if a contract is offered and the agent shirks, then we have $\bar{W} = c + \delta\bar{W}$, giving $\bar{W} = \frac{c}{1-\delta}$.

allowing us to write

$$\begin{aligned}\bar{W} &= \max_{s,W} p\pi(1-s) + \delta(1-p)W \\ \text{s.t. } \bar{W} &\geq W \\ p\pi s - c + \delta(1-p) &\left[\frac{p\pi - c}{1 - \delta(1-p)} - W \right] = \frac{p\pi - 2c - \delta p \bar{W}}{1 - \delta(1-p)}.\end{aligned}$$

Now notice that the objective and the final constraint involve identical linear tradeoffs of s versus W . We can thus assume that $W = \bar{W}$, allowing us to write the problem as

$$\bar{W} = \max_s p\pi(1-s) + \delta(1-p)\bar{W} \quad (87)$$

$$\text{s.t. } p\pi s - c + \delta(1-p) \left[\frac{p\pi - c}{1 - \delta(1-p)} - \bar{W} \right] = \frac{p\pi - 2c - \delta p \bar{W}}{1 - \delta(1-p)}. \quad (88)$$

We now show that this implies $\bar{W} = c/(1 - \delta)$. From (87), we have (letting s^* be the maximizer, subtracting c from both sides, and rearranging)

$$p\pi s^* - c = p\pi + \delta(1-p)\bar{W} - \bar{W} - c.$$

Now using (88), we can write this as

$$\frac{p\pi - 2c - \delta p \bar{W}}{1 - \delta(1-p)} - \delta(1-p) \left[\frac{p\pi - c}{1 - \delta(1-p)} - \bar{W} \right] = p\pi + \delta(1-p)\bar{W} - \bar{W} - c,$$

or, isolating \bar{W} ,

$$\bar{W} \left[\frac{\delta p}{1 - \delta(1-p)} - 1 \right] = \frac{p\pi - 2c}{1 - \delta(1-p)} - \delta(1-p) \frac{p\pi - c}{1 - \delta(1-p)} - [p\pi - c],$$

or (simplifying the left side and multiplying by -1),

$$\frac{(1 - \delta)\bar{W}}{1 - \delta(1-p)} = (p\pi - c) + \frac{\delta(1-p)(p\pi - c)}{1 - \delta(1-p)} - \frac{p\pi - 2c}{1 - \delta(1-p)},$$

or

$$\bar{W} = \frac{[1 - \delta(1-p)](p\pi - c) + \delta(1-p)(p\pi - c) - (p\pi - 2c)}{1 - \delta} = \frac{c}{1 - \delta}.$$

E.5.2 Proof of Lemma 20

We consider an artificial game in which the principal is free of sequential rationality constraints. The principal names, at the beginning of the game, a pair of sequence $\{t_n\}_{n=0}^{\infty}$

and $\{s_n\}_{n=0}^\infty$ such that, barring a success, the principal makes an offer s_n at time t_n . To preserve feasibility, we must have $t_{n+1} - t_n \geq \Delta$, with strict inequality if there is delay. The principal's objective is to minimize the agent's payoff subject to the constraints that the agent be willing to exert effort in response to any offer, and that the principal's payoff in the continuation game starting at each period is at least V^* . We show that the bounds on the agent's payoff given by $\frac{c}{1-\delta}$ (if $\Psi^\dagger < 1$) and $\frac{V^*}{\delta p}$ (if $\Psi^\dagger > 1$) apply to this artificial game. The bounds must then also hold in the original game. Since we have equilibria of the original game whose payoffs approach (as $\Delta \rightarrow 0$) the proposed payoff in each case, this establishes the result.

First, we note that $t_0 = 0$, since otherwise the principal could increase her payoff by eliminating the initial delay without compromising the constraints. Next, each offer s_n must cause the agent's incentive constraint to bind. Suppose to the contrary that at some time t_n the agent's incentive constraint holds with strict inequality. Then replacing the offer s_n with the (larger) value s_n^* that causes the agent's constraint to bind, while leaving continuation play unaffected, preserves the agent's incentives (since the continuation value of every previous period is decreased, this only strengthens the incentives in previous periods) while increasing the principal's and reducing the agent's payoff, a contradiction.

Let \underline{W} be the agent's minimum equilibrium payoff. Because the agent's incentive constraint always binds, \underline{W} must equal the expected payoff from persistent shirking, and hence is given by

$$\underline{W} = c \sum_{n=0}^{\infty} e^{-\delta t_n}. \quad (89)$$

Notice that the continuation payoff faced by the agent at each time t_n must be at least \underline{W} , since otherwise \underline{W} is not the lowest equilibrium payoff possible for the agent. Next, we claim that each such continuation payoff equals \underline{W} . If this is not the case for some t_n , then we can construct an alternative equilibrium featuring the same sequence of times and offers for $n = \{0, \dots, t_n - 1\}$, and then continuing with an equilibrium in the resulting continuation game that gives payoff \underline{W} . Because the continuation value at time t_n has been reduced, this allows us to reduce the first-period value s_0 while still preserving all of the agent's incentive constraints. The resulting lower first-period payoff and lower continuation value decrease the agent's payoff (and increase the principal's), a contradiction.

Using (89), this in turn implies that

$$\begin{aligned} \frac{\underline{W}}{c} &= \sum_{n=0}^{\infty} e^{-\delta t_n} \\ &= 1 + \sum_{n=1}^{\infty} e^{-\delta t_n} \\ &= 1 + e^{-\delta t_1} \sum_{n=1}^{\infty} e^{-\delta(t_n - t_1)} \\ &= 1 + e^{-\delta t_1} \frac{\underline{W}}{c}, \end{aligned}$$

where the final equality uses the fact that the agent's continuation value at t_1 is \underline{W} . We can repeat this exercise from the point of view of time t_1 , giving

$$\begin{aligned}\frac{\underline{W}}{c} &= \sum_{n=1}^{\infty} e^{-\delta(t_n-t_1)} \\ &= 1 + e^{-\delta(t_2-t_1)} \sum_{n=2}^{\infty} e^{-\delta(t_n-t_2)} \\ &= 1 + e^{-\delta(t_2-t_1)} \frac{\underline{W}}{c}.\end{aligned}$$

We can conclude in this fashion, concluding that there exists some Ψ such that for all $n \geq 1$,

$$t_n - t_{n-1} = \Psi \Delta.$$

However, we have characterized the equilibria that feature such a constant value of Ψ , finding that the only such equilibrium gives payoff $W^* = \frac{c}{1-\delta}$ when $\Psi^\dagger < 1$ and that the agent's lowest payoff from such an equilibrium is $\frac{V^*}{\delta p}$ if $\Psi^\dagger > 1$.

E.5.3 Proof of Lemma 21

We consider an equilibrium with payoffs (W_0, V_0) . We are interested in an upper bound on the ratio $\frac{V_0 - V}{W_0 - \underline{W}}$, which we denote by ζ . It suffices to consider an equilibrium in which a period-0 mixture with probability $(1 - x_0)$ prompts the players to continue with equilibrium payoffs $(\underline{W}, \underline{V})$, and with probability x_0 calls for a current contract s , followed by a period-1 mixture attaching probability $1 - x_1$ between continuation payoffs $(\underline{W}, \underline{V})$ and probability x_1 to continuation play with payoffs (W_1, V_1) , and so on. In addition, we can assume that any contract offered to the agent induces the agent to work.⁴⁵ Hence,

⁴⁵Any such contract is part of an extreme equilibrium. Suppose we have a contract that does not induce effort, and hence gives payoffs $-c + \delta V$ and $c + \delta W$ to the principal and agent, respectively, for some continuation payoffs (W, V) . There exists an alternative equilibrium with the same continuation payoffs, but in which the principal induces effort by offering a share s satisfying

$$c + \delta W = (1 - s)p\pi + \delta(1 - p)W.$$

Solving this expression gives $p\pi - c - \delta pW = p\pi s$, and hence a principal payoff of $p\pi - 2c - \delta pW + \delta(1 - p)V$. It is then a contradiction to our hypothesis that we are dealing with an extreme equilibrium, hence establishing the result, to show that this latter payoff exceeds $-c + \delta V$, or $p\pi - 2c - \delta pW + \delta(1 - p)V > -c + \delta V$, which is $p\pi - c > \delta p(W + V)$, or

$$\frac{p\pi - c}{\delta p} > V + W.$$

The left side is an upper bound on the value of the project without an agency problem, giving the result.

we have

$$\begin{aligned} V_0 &= x_0 [p\pi s - c + \delta(1-p)[x_1 V_1 + (1-x_1)\underline{V}]] + (1-x_0)\underline{V} \\ W_0 &= x_0 [p\pi(1-s) + \delta(1-p)[x_1 W_1 + (1-x_1)\underline{W}]] + (1-x_0)\underline{W} \\ &\geq x_0 [c + \delta[x_1 W_1 + (1-x_1)\underline{W}]] + (1-x_0)\underline{W}, \end{aligned}$$

where the inequality is the agent's incentive constraint. Setting an equality in the incentive constraint, we can solve for

$$p\pi s = p\pi - c - \delta p[x_1 W_1 + (1-x_1)\underline{W}].$$

Using this to eliminate the share s from the principal's payoff, and returning to the agent's binding incentive constraint, we obtain

$$\begin{aligned} V_0 - \underline{V} &= x_0 [p\pi - 2c - \delta p[x_1 W_1 + (1-x_1)\underline{W}] + \delta(1-p)[x_1 V_1 + (1-x_1)\underline{V}] - \underline{V}] \\ W_0 - \underline{W} &= x_0 [c + \delta[x_1 W_1 + (1-x_1)\underline{W}] - \underline{W}], \end{aligned}$$

and hence

$$\zeta := \frac{V_0 - \underline{V}}{W_0 - \underline{W}} = \frac{p\pi - 2c - \delta p[x_1(W_1 - \underline{W}) + \underline{W}] + \delta(1-p)[x_1(V_1 - \underline{V}) + \underline{V}] - \underline{V}}{c + \delta[x_1(W_1 - \underline{W}) + \underline{W}] - \underline{W}}.$$

We obtain an upper bound on this expression by first taking $V_1 - \underline{V} = \zeta(W_1 - \underline{W})$ on the right side and then rearranging to obtain

$$\zeta \leq \frac{p\pi - 2c - \delta p[x_1(W_1 - \underline{W}) + \underline{W}] + (1-\delta)\underline{V}}{c + \delta[x_1 p(W_1 - \underline{W}) + \underline{W}] - \underline{W}}.$$

We now note that $W_1 - \underline{W}$ appears negatively in the numerator and positively in the denominator, so that an upper bound on ζ is obtained by setting $W_1 - \underline{W} = 0$ on the right side, giving

$$\zeta \leq \frac{p\pi - 2c - \delta p\underline{W} - (1-\delta(1-p))\underline{V}}{c - (1-\delta)\underline{W}} = \frac{\delta p}{1-\delta}, \quad (90)$$

where the final equality is obtained by using $\underline{W} = \frac{1-\delta}{\delta p}\underline{V}$ to eliminate \underline{W} , and then simplifying.

E.5.4 Proof of Lemma 22

We assume that $x_0, x_1 \in (0, 1)$ and establish a contradiction. Using the incentive constraint, we can write

$$\begin{aligned} W_0 &= x_0 [c + \delta[x_1 W_1 + (1-x_1)\underline{W}]] + (1-x_0)\underline{W} \\ V_0 &= x_0 [p\pi - 2c - \delta p[x_1 W_1 + (1-x_1)\underline{W}] + \delta(1-p)[x_1 V_1 + (1-x_1)\underline{V}] - (1-x_0)\underline{V}]. \end{aligned}$$

We now identify the rates at which we could decrease x_1 and increase x_0 while preserving the value W_0 . Thinking of x_0 as a function of x_1 , we can take a derivative of this expression for W_0 to find

$$\frac{dW_0}{dx_1} = \frac{dx_0}{dx_1} \frac{W_0 - \underline{W}}{x_0} + \delta x_0 (W_1 - \underline{W}) = 0,$$

and then solve for

$$\frac{dx_0}{dx_1} = \delta x_0^2 \frac{W_1 - \underline{W}}{W_0 - \underline{W}}.$$

Now let us differentiate V_0 to find to find

$$\begin{aligned} \frac{dV_0}{dx_1} &= \frac{dx_0}{dx_1} \frac{V_0 - \underline{V}}{x_0} + \delta x_0 [(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})] \\ &= -\delta x_0 \frac{W_1 - \underline{W}}{W_0 - \underline{W}} (V_0 - \underline{V}) + \delta x_0 [(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})]. \end{aligned}$$

It is a contradiction to show that this derivative is negative, since then we could increase the principal's payoff, while preserving the agent's by decreasing x_1 . Eliminating the term δx_0 and multiplying by $W_0 - \underline{W} > 0$, we have

$$[(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})](W_0 - \underline{W}) - (V_0 - \underline{V})(W_1 - \underline{W}) \leq 0.$$

We now substitute for $W_0 - \underline{W}$ and $V_0 - \underline{V}$ to obtain

$$\begin{aligned} & [(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})] x_0 [c + \delta[x_1 W_1 + (1-x_1)\underline{W}] - \underline{W}] \\ & - x_0 [p\pi - 2c - \delta p[x_1 W_1' + (1-x_1)\underline{W}] + \delta(1-p)[x_1 V_1' + (1-x_1)\underline{V}] - \underline{V}] (W_1 - \underline{W}) \\ & \leq 0. \end{aligned}$$

Deleting the common factor x_0 and canceling terms, this is

$$\begin{aligned} & [(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})] [c - (1-\delta)\underline{W}] \\ & - [p\pi - 2c - \delta p\underline{W} + \delta(1-p)\underline{V} - \underline{V}] (W_1 - \underline{W}) \leq 0. \end{aligned}$$

Rearranging, we have

$$\frac{(1-p)(V_1 - \underline{V}) - p(W_1 - \underline{W})}{W_1 - \underline{W}} \leq \frac{p\pi - 2c - \delta p\underline{W} - (1-\delta(1-p))\underline{V}}{c - (1-\delta)\underline{W}},$$

which follows immediately from the inequality in (90) from the proof of Lemma 21.

E.5.5 Proof of Proposition 9

The two differential equations (85)–(86) have as solutions, for some $C_1, C_2 \in \mathbb{R}$,

$$w(t) = \frac{c}{r} + C_1 e^{rt}, \text{ and } v(t) = \frac{\psi - \sigma}{\sigma + 1} \frac{c}{r} - C_1 e^{rt} + (C_1 + C_2) e^{r(1+\sigma)t}.$$

If $0 < \psi < \sigma$ (the case of a patient project), then, since the first term of the principal's payoff is strictly negative, it must be that either C_1 , or $C_1 + C_2$ is nonzero. Since the solution must be bounded, this implies, as expected, that effort cannot be supported (without delay) indefinitely. If effort stops at time T , then, since $w(T) = 0$, $C_1 e^{rT} = -c/r$, and C_2 is then obtained from $v(T) = 0$. Eliminating T then yields the following relationship between $v(0)$ and $w(0)$, written simply as v and w :

$$v = \frac{\psi + 1}{\sigma + 1} \left[1 - \left(1 - \frac{rw}{c} \right)^{\sigma+1} \right] \frac{c}{r} - w.$$

We let w^\dagger denote the unique strictly positive root of the previous expression. If $w \in [0, w^\dagger]$, then $v \geq 0$, and these are the values that can be obtained for times T for which the principal's payoff is positive. This yields the result for patient projects. For reference, the Markov equilibrium in this region is given by $(w, v) = (\frac{\psi c}{\sigma r}, 0)$.

Now consider impatient projects, or $\psi > \sigma$, so that the principal's payoff in the stationary full-effort equilibrium is positive. We need to describe the equilibrium payoffs of potential stationary-outcome equilibria with delay. We encompass delay in the discount rate. That is, players discount future payoffs at rate $r\lambda$, for $\lambda \geq 1$. The payoffs to the agent and principal, under such a constant rate, are

$$w = \frac{c}{r\lambda}, \quad v = \frac{\lambda\psi - \sigma}{\sigma + \lambda} \frac{c}{r\lambda}.$$

There exists at most one value of $\lambda > 1$ for which the principal's payoff is equal to that obtained for $\lambda = 1$, namely

$$\lambda = \frac{\sigma(\sigma + 1)}{\psi - \sigma},$$

which is larger than one if and only if $\psi < \sigma(\sigma + 2)$. As before, if $\psi > \sigma(\sigma + 2)$, then we have the case of a very impatient project, for which there is no other equilibrium payoff than the Markov payoff (w^*, v^*) .

Let us then focus on moderately patient projects for which

$$\psi \in (\sigma, \sigma(\sigma + 2)),$$

in which case $\lambda > 1$, so that there exists an equilibrium in which constant funding is provided, but at a slower rate than possible. The agent's payoff in this equilibrium is

$$\underline{w} = \frac{c}{r\lambda} = \frac{\psi - \sigma}{\sigma(\sigma + 1)} \frac{c}{r}.$$

We may now solve the differential equations with boundary condition $v(T) = v^*$, $w(T) = \underline{w}$ for an arbitrary $T \geq 0$. Eliminating T gives the following relationship between $v = v(0)$ and $w = w(0)$:

$$v = \frac{c}{r} \left(\frac{\psi + 1}{\sigma + 1} + \frac{\psi - \psi\sigma - 2\sigma}{\psi - \sigma^2 - 2\sigma} \left(\frac{\psi - \sigma^2}{\sigma^2 + \sigma} \right)^{-\sigma} \left(1 - \frac{rw}{c} \right)^{\sigma+1} \frac{c}{r} \right) - w,$$

completing the results for moderately impatient projects.

F Online Appendix: Proofs, Commitment

F.1 Proof of Proposition 6

Given $w(q)$, we would like to characterize the value of $v(q)$ that maximizes the principal's payoff among equilibrium payoffs $\{w(q), v(q)\}$. We simplify the argument by assuming the players have access to a public randomization device, describing at the end of the proof the modifications required if no such device is available. We make use of the fact that we have a “worst continuation” that simultaneously delivers the worst possible equilibrium payoffs to both the principal and the agent. In particular, given that the principal has commitment power, this worst continuation is a full-stop continuation that delivers payoff zero to both agents.

The proof is quite similar in structure to the corresponding proof of Proposition 2 for the case without commitment (cf. Section C.2). Fix a value q . Then, because the worst equilibrium gives payoffs $(0, 0)$, it is immediate that any equilibrium payoff $(w(q), v(q))$ can be achieved as a mixture of an equilibrium that features no delay and the worst equilibrium. Hence, we can restrict attention to sequences $(x(q), s(q))$, where, given posterior q , full-stop continuation is played with probability $1 - x(q)$ (determined by a public randomization device); and if not, a share $s(q)$ is offered in that period that induces the agent to work.

F.1.1 A Preliminary Inequality

We fix a posterior probability and let $w(q)$ and $v(q)$ be equilibrium values. We argue that

$$\frac{v(q)}{w(q)} \leq \frac{\pi - 2c}{c} := \zeta.$$

In particular, any equilibrium exertion of effort on the part of the agent creates a discounted surplus of $(pq\pi - c)\Delta$, where the discounting reflects the delay until the effort is exerted and the probability that the interaction may terminate before reaching such effort. Of this surplus, at least $c\Delta$ must go to the agent, since otherwise the agent's

incentive constraint is surely violated. The ratio of principal to agent payoffs can then never exceed $(p\pi - 2c)/c$.

F.1.2 Front-Loading Effort

Fix a posterior q . We first note that

$$\begin{aligned} v(q) &= x(q) [(pqs\pi - c)\Delta + \delta(\Delta)(1 - pq\Delta)[x(\varphi(q))v(\varphi(q))]] \\ w(q) &= x(q) [pq(1 - s)\pi\Delta + \delta(\Delta)(1 - pq\Delta)[x(\varphi(q))w(\varphi(q))]] \\ &\geq x(q) [c\Delta + \delta(\Delta)[x(\varphi(q))\theta w(\varphi(q))], \end{aligned}$$

where $\varphi(q)$ is the posterior belief obtained from q given a failure (cf. (31)). The inequality is the agent's incentive constraint and $\theta > 1$ is given by

$$\theta = \frac{q}{\varphi(q)} = \frac{1 - p\Delta q}{1 - p\Delta},$$

and hence is the ratio of the current posterior to next period's posterior, given a failure. We have used here the fact that the continuation values, relevant for posterior $\varphi(q)$, can be written as convex combinations of equilibrium payoffs $(w(\varphi(q)), v(\varphi(q)))$ and the full-stop continuation values $(0, 0)$. Note that in writing this convex combination, we take $w(\varphi(q))$ and $v(\varphi(q))$ to be the interim values, i.e., values at the point at which the posterior is $\varphi(q)$ and precisely Δ time has elapsed since the previous offer. The equilibrium generating these values may yet entail some delay.

Let us simplify the notation by letting $x(q) = x$, $v(q) = v$, $w(q) = w$, $x(\varphi(q)) = \tilde{x}$, $v(\varphi(q)) = \tilde{v}$, $w(\varphi(q)) = \tilde{w}$, and let us drop the explicit representation of Δ . Setting an equality in the agent's incentive constraint and rearranging gives

$$pqs\pi = (pq\pi - c) + \delta(1 - pq)\tilde{x}\tilde{w} - \delta\tilde{x}\theta\tilde{w}.$$

Using this to eliminate the variable s from the value functions gives

$$v = x [(pq\pi - 2c) + \delta(1 - pq)\tilde{x}\tilde{w} - \delta\tilde{x}\theta\tilde{w} + \delta(1 - pq)\tilde{x}\tilde{v}], \quad (91)$$

$$w = x [c + \delta\tilde{x}\theta\tilde{w}]. \quad (92)$$

We now show that it is impossible for x and \tilde{x} to both be interior. Suppose they are. Then we consider an increase in x and an accompanying decrease in \tilde{x} , effectively moving effort forward. We keep w constant in the process, and show that the result is to increase v , a contradiction.

First, we fix the constraint by differentiating (92) to find

$$\frac{dw}{d\tilde{x}} = \frac{dx}{d\tilde{x}} \frac{w}{x} + \delta x \theta \tilde{w},$$

and hence, setting $\frac{dw}{d\tilde{x}} = 0$,

$$\frac{dx}{d\tilde{x}} = -\delta x^2 \theta \frac{\tilde{w}}{w}. \quad (93)$$

Differentiating (91) and using (93), we have

$$\begin{aligned} \frac{dv}{d\tilde{x}} &= \frac{dx}{d\tilde{x}} \frac{v}{x} + \delta x ((1 - pq - \theta)\tilde{w} - (1 - pq)\tilde{v}) \\ &= -\delta x \theta \frac{\tilde{w}}{w}(v) + \delta x ((1 - pq - \theta)\tilde{w} - (1 - pq)\tilde{v}). \end{aligned}$$

It concludes the argument to show that this derivative is negative. Multiplying by w , the requisite inequality is

$$w((1 - pq - \theta)\tilde{w} + (1 - pq)\tilde{v}) - \tilde{w}v\theta < 0.$$

Substituting for v and w from (69)–(70) and dropping the common factor x , this is

$$\begin{aligned} &[(1 - pq - \theta)\tilde{w} + (1 - pq)\tilde{v}](c + \delta\tilde{x}\theta\tilde{w}) \\ &< \tilde{w}\theta[(1 - \delta)(p\pi - 2c) + \delta(1 - pq - \theta)\tilde{x}\tilde{w} \\ &\quad + \delta(1 - pq)\tilde{x}\tilde{v}]. \end{aligned}$$

We can then note that the terms involving \tilde{x} cancel, at which point the expression simplifies to

$$(1 - p - \theta) + (1 - p)\frac{\tilde{v}}{\tilde{w}} < \theta \frac{(p\pi - 2c)}{c},$$

for which, using the definition of ζ , it suffices that $(1 - p - \theta) + (1 - p)\zeta < \theta\zeta$, which is immediate.

An implication of this result is that $x(q)$ is interior for at most one value of q . This in turn implies that the public randomization device is required only for one value of q . If no public randomization device is available, we can construct an equilibrium in which no values of $x(q)$ are interior that approximates the equilibrium examined here. As $\Delta \rightarrow 0$, the public randomization device becomes unimportant and the approximation becomes arbitrarily sharp.