

Essays on Dynamic Games and Mechanism Design


by
Ruitian Lang

B.S. Mathematics and Physics
Massachusetts Institute of Technology, 2009

SUBMITTED TO THE DEPARTMENT OF ECONOMICS IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY IN ECONOMICS
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

©2014 Ruitian Lang. All Rights Reserved.



The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature redacted

Author

Department of Economics

April 15, 2014

Certified by

Signature redacted

Daron Acemoglu

Elizabeth and James Killian Professor of Economics

Certified by

Signature redacted

Robert Gibbons

Sloan Distinguished Professor of Management and Economics

Accepted by

Signature redacted

Michael Greenstone

3M Professor of Environmental Economics, Public Economics, Labor Economics
Chairman, Department Committee of Graduate Thesis

Essays on Dynamic Games and Mechanism Design

by
Ruitian Lang

Submitted to the Department of Economics on June 8th, 2014 in Partial
Fulfillment of the Requirements for the Degree of Doctor of Philosophy in
Economics

ABSTRACT

The dissertation considers three topics in dynamic games and mechanism design. In both problems, asymmetric information causes inefficiency in production and allocation. The first chapter considers the inefficiency from the principal's inability to observe the agent's effort or cost of effort, and explores its implication to the principal's response to the combination of the output and the signal about the cost of effort. For example, the principal may punish the agent more harshly for low output when signals suggest that cost of effort is high when the effort is of high value for the principal. This chapter also classifies the long-run behavior of the relationship between the principal and the agent. Depending on whether the agent is strictly risk-averse and whether he is protected by limited liability, the state of the relationship may or may not converge to a stationary state and the stationary state may or may not depend on the initial condition.

The second chapter considers the re-allocation of assets among entrepreneurs with different matching qualities, which contributes to the growth of the whole economy. Due to reasons that are not explicitly modeled, assets are not automatically allocated to entrepreneurs who are best at operating them from the beginning, and this inefficiency is combined with inefficiency in the asset market and potential imperfection of labor contracting. When asset re-allocation can become a main source of economic growth, this chapter argues that imperfection in the labor contracting environment may boost the economic growth.

The third chapter assumes that the agent's output is contractible but he can privately acquire more information about his cost of production prior to contracting. Compared to the optimal screening contract, the principal's contract in this case must not only induce the agent to "tell the truth", but also to give the agent the incentive to acquire appropriate amount of information. This may create distortion of allocation to the most

efficient type and whether this happens is related to the marginal loss incurred by the principal from the cost of information acquisition.

Thesis Supervisor: Daron Acemoglu

Title: Elizabeth and James Killian Professor of Economics

Thesis Supervisor: Robert Gibbons

Title: Sloan Distinguished Professor of Management and Economics

Thesis Supervisor: Juuso Toikka

Title: Assistant Professor of Economics

Contents

Contents	5
List of Figures	7
List of Tables	9
1 Optimal Informal Incentives and Long-Run Dynamics in a Fluctuating Environment	1
1.1 Introduction	1
1.1.1 Related Literature	7
1.2 The Model	9
1.3 Preliminaries	12
1.3.1 The equilibrium frontier	12
1.3.2 The first-order approach	16
1.4 Optimal informal incentives	20
1.4.1 Optimal response to the fluctuating environment	20
1.4.2 Evolution of effort level	24
1.5 Long-run Dynamics	27
1.5.1 The main result	27
1.5.2 Sketch of the proof	34
1.5.3 Long-run effects of the initial state	37
1.6 Extension	41
1.7 Conclusion	43
1.8 Appendices	45
1.8.1 Appendix A: proofs	45
1.8.2 Appendix B: long-run dynamics with commitment	63
2 Relationships in a Market: Contract Enforcement and Asset Re-allocation	69
2.1 Introduction	69
2.2 The Model	72
2.3 The Partnership Game and Determination of Asset Price	76
2.3.1 The strong contracting environment	77
2.3.2 The weak contracting environment	80
2.4 The Equilibrium Growth Path	84
2.4.1 Comparative statics	86
2.4.2 A "two-type" example	88
2.5 A Continuum of Contracting Environment	92
2.6 Conclusion	95

2.7	Appendix	97
3	Optimal Mechanisms with Information Acquisition	109
3.1	Introduction	109
3.2	The Model	112
3.3	The Two-Type Case	115
3.3.1	Additional distortion and inefficiency at the top	115
3.3.2	Investing in a risky project: an example	118
3.4	The Linear Model	122
3.4.1	Value of information and the role of the prior	122
3.4.2	Characterizing the conditional optimal allocations	125
3.4.3	The discrete type space	126
3.5	Applications	130
3.5.1	Binary information acquisition	130
3.5.2	A success-failure experiment	137
3.6	Conclusion	140
3.7	Appendix	143
4	References	149

List of Figures

1.1	The equilibrium frontier under limited liability	14
1.2	Branches of the frontier under limited liability	14
1.3	The equilibrium frontier without limited liability	15
1.4	Branches of frontier without limited liability	16
1.5	Evolution of the agent's effort in the bad environment	33
1.6	Simulated cumulative distribution function of the invariant distribution π	34
2.1	Fraction of low-quality asset owners who sell their assets.	91
2.2	Long-run aggregate output as a function of a_L	91

List of Tables

Chapter 1

Optimal Informal Incentives and Long-Run Dynamics in a Fluctuating Environment

1.1 Introduction

In a long-run relationship between a risk-neutral principal and a risk-averse agent, the scope of their cooperation can change over time. In its most extreme form, the change involves the termination of their relationship: the principal fires the agent. Even if the relationship continues forever, the agent's effort level can also change over time depending on realizations of outputs on the equilibrium path. Two questions can be asked about the dynamics of the relationship. First, how does the state of the relationship changes in response to the agent's effort? Secondly, will the relationship converges to a stationary state in the long run, and if so, does the stationary state depend on its initial state?

The two questions are related. The answer to the first question characterizes the transition from the state in one period to the state in the next. The long-run limit can be thought of as a result of iterating this transition law, although the formal analysis is much more complicated as the state can take values in a continuum. The characterization of the optimal response to the agent's output is a solved problem under the standard assumptions of the moral hazard theory (such as Monotone Likelihood Ratio

Property and Convex Distribution Function Condition). However, this framework is not general enough to analyze some relationships within and between firms and explain the observed empirical patterns. The common feature of the applications briefly described below has the common feature that there are shocks to the agent's cost of effort, and it is natural to study the optimal response to public signals about the shock, also referred to as the environment, in addition to outputs.

The model studied in this paper can be used to study at least three types of relationships in the real world. In a relationship between a buyer and a seller, the seller has more information about cost of production which may fluctuate over time. While the buyer does not know the seller's cost precisely, she may have access to news that is correlated with the seller's cost. In the problem of corporate governance, the CEO is expected to make effort to generate profits and to resist temptation of diverting the firm's cash. The CEO has more information about the effort required to reduce cost or to find new investment opportunities and the temptation of diverting than the board, but the board can use competitors' performances as proxies to the environment. Finally, the model can be used to study the employment relationship in which the employer can choose how much bonus to pay to the agent but her cost of cash payment depends on the liquidity of the firm which is subject to exogenous shocks.

Using a data set of relationships between Kenyan flower exporters and foreign buyers, Macchiavello and Morjaria (2012) shows that exporters performed differently during a conflict in Kenya that substantially increased their costs of labor. In addition, dividing exporters into a group directly affected by the conflict and a group not directly affected, the authors find that an exporter's contract in the following season (measured by its existence, price, and quantity) is more responsive to his performance during the conflict period if he is in the group directly affected by the conflict. This finding is hard to explain by both classical theories which imply that the principal (foreign buyer) should insure the agent (exporter) against low output generated by bad luck and theories of relative performance evaluation¹. The study of optimal informal incentives in this paper explains this finding and the study of long-run dynamics asks further how the performance differences among exporters arise and whether they should exist in the long run.

¹See Jenter and Kanaan (forthcoming) for a theoretical framework of relative performance evaluation and empirical evidence from corporate governance that suggests that such evaluation is only used to a very limited degree.

We model the evolution of state in an optimal equilibrium of the repeated game between the principal and the agent as a Markov process, and the state is their continuation payoffs. The study of the transition probabilities characterizes the optimal response to the output and the environment, or the short-run dynamics. On the other hand, the ergodic theory of such processes can be used to study the long-run performance of the relationship for any given discount factor. As alluded to before, unlike Markov chains on a countable state space, this Markov process has a continuum state space, and its ergodic theory requires some regularity conditions. An important step is to establish the connection between these regularity conditions and assumptions on the principal-agent relationship. For tractability and interpretation reasons, the baseline model focuses on a simple performance problem: whether the performance of the relationship is robust to the shock to the agent's cost. Then it is shown that the ergodic theory can be applied to a more general setting.

There are at least three possible types of long-run behaviors of the optimal equilibrium. The state of the relationship may converge to a unique stationary state independent of its initial state, may converge to a unique wide-spread distribution independent of its initial state, or may converge to a limit that depends on the initial state. It will be shown that in the baseline setting, the limit distribution is independent of the initial state and the path, and is non-degenerate: its support is all the extreme points of the equilibrium payoff set on the Pareto frontier. It will also be shown how the other two types of limit behaviors arise when we change assumptions on the agent's payoff function.

These results have implications for the long-run performance difference problem in organizational economics (see Chassang (2010) for example)². Due to the dependence of the agent's effort level on history, organizations with the same production function and the same initial condition can develop different expected outputs over time. When the limit distribution of state is non-degenerate, performance differences exist in the long run if they exist in the short run. On the other hand, when the limit distribution is independent of the initial state, the organization that currently enjoys the highest performance will not lead forever: the performance leader in the group changes over time. This mean-reversion behavior is due to the agent's diminishing marginal utility

²See Bailey, Hulten, Campbell, Bresnahan, and Caves (1992), Syverson (2004), Foster, Haltiwanger, and Syverson (2008) for empirical studies of persistent performance differences, and see Syverson (2011) for a survey. Gibbons and Henderson (2012) propose that some of these performance differences are rooted in management practices that rely on relational contracts.

of consumption, which makes it hard to provide him with incentives when his expected continuation payoff is high due to past successes.

By focusing on temporary fluctuations in a relationship where the principal cannot commit to a long-term contract, the paper leaves out dynamics driven by (a) Bayesian updating about the value of the relationship and (b) a complete long-term contract. The value of a relationship may be affected by the agent's ability and the profitability of the principal's business, and Bayesian updating about the value is a potentially important driver of dynamics, but it is hard to attribute observed performance differences among well established firms or plants in mature industries to Bayesian updating. In contrast, temporary fluctuations always exist in these firms or plants and generate performance difference in the long run. Therefore, the analysis in this paper complements the analysis of dynamics driven by Bayesian updating.

Dynamics under the optimal long-term contract have been studied fruitfully in the literature, especially in continuous time. However, a complete long-term contract between the principal and the agent does not always exist. For example, the principal can freely fire the agent or refuse to renew a short-term contract in many applications. It will be shown that the long-run limit of the dynamics under a complete contract is very different from the limit without commitment. In particular, under the optimal complete contract, the relationship eventually converges to a stationary state in which the agent receives payment in every period but does not exert any effort, which can never happen when the principal cannot commit to such a contract.

This paper models the relationship between a principal and an agent as a repeated game: no party has commitment power. The risk-neutral principal can reward the risk-averse agent with cash, and the agent cannot save, so the principal controls the agent's consumption path. On the other hand, the agent bases his choice of effort on a privately observed exogenous state of environment, which is either "good" or "bad". The agent's cost of effort is higher in the bad environment. Output is a noisy measure of the agent's effort. At the end of each period, a signal about the environment is publicly observed. Limited liability is often an important constraint on the principal's payment to the agent, so the variant model with limited liability will be considered in parallel to the baseline model without limited liability.

There are three main sets of results. The first is concerned with how dynamics arise in the first place. As long as the agent is expected to exert effort in the good environment, his continuation payoff increases if he produces high output and decreases if he produces low output. This result is independent of the signal about the environment. On the other hand, the magnitude of the difference in the agent's continuation payoff depends on that signal. In turn, the agent's continuation payoff determines his effort in the next period, and this feedback drives the dynamics of the relationship. When the common discount factor of the principal and the agent on future payoffs is sufficiently high, the agent's efforts in both environments decrease with his continuation payoff. Therefore, an agent who has accumulated many high outputs tends to exert lower effort in the future until his continuation drops to a lower level as a result of the lower output. However, when the agent is protected by limited liability, the result is very different. When the agent's expected payoff is very low, it is very costly or even impossible for the principal to provide the agent with incentives, so the agent stops exerting effort as his expected payoff approaches zero. Therefore, the performance of a principal-agent relationship without limited liability exhibits mean-reversion behavior, while the performance of one with limited liability does not always do so.

The second set of results provides sufficient conditions under which the state of the relationship converges over time in the optimal equilibrium. When the agent is strictly risk-averse, and the distributions of output and the signal about the environment satisfy some regularity conditions, the optimal continuation equilibrium is uniquely determined by the principal's and the agent's expected payoffs, and the probability distribution of their payoffs converge in the long run. Moreover, the limit distribution is independent of the initial state and thus of any information in a finite horizon. In the baseline model without limited liability, the support of the limit distribution is all the extreme points on the Pareto frontier of the equilibrium payoff set. If this distribution is thought of as the distribution of states of independent organizations, the fact that it has a broad support means that there will be performance difference in the long run, but the top and bottom performers will be re-shuffled over time. When the agent is protected by limited liability, however, the limit distribution is the unit mass at the payoff of the static equilibrium; in other words, eventually the principal and the agent start playing the static equilibrium forever. Intuitively, the agent stops exerting effort when his payoff is low under limited liability, which implies that the relationship cannot bring itself out

of a low state, and thus will be eventually absorbed by the lowest possible state, the static equilibrium payoff.

As the above result only holds when the agent is strictly risk-averse, the third result illustrates what happens when the agent is risk-neutral at sufficiently high consumption levels and is protected by limited liability. In this case, every optimal equilibrium path converges to a stationary state, but the limit state depends on the initial state. Therefore, relationships with the same production function but different initial states may have very different expected output in the long run. In particular, some relationships will enter a stationary regime with efficient effort from the agent in every period, while other relationships will be terminated. This shows that persistent performance difference may arise purely from the optimal dynamic response to the agent's output and information about a fluctuating environment.

The main contribution of this paper is to characterize the optimal equilibrium of a principal-agent relationship without commitment for all discount factors, and to characterize different types of long-run behaviors of the optimal equilibrium in a unified framework. Many papers in the dynamic game literature are content with the characterization of the equilibrium payoff set and the difference equation that governs the equilibrium dynamics. However, in a continuum state space it is not obvious how one can iterate this difference equation to learn about the long-run behavior of the relationship. Sometimes the equilibrium dynamics is simple enough and the long-run behavior of the relationship is easy to predict. (See Li and Matouschek (2013) for an example.) The current paper provides tools that can potentially be used to characterize the long-run behavior of more complicated relationships.

The remainder of the paper is organized as follows. Section 1.1 reviews the related literature. Section 2 presents the model. Section 3 characterizes the equilibrium payoff set and derives the difference equation that describes the evolution of the agent's expected continuation payoff. Section 4 applies results in Section 3 to the study of short-run dynamics. Section 5 states and proves the main result on the long-run dynamics, and also gives an example that violates some assumptions of the main convergence theorem and generates persistent performance differences. Section 6 applies the technique used in Section 5 to a setting that allows for continuous effort and more general monitoring structure. Section 7 concludes.

1.1.1 Related Literature

The paper is related to several literatures. First, it uses the framework of repeated games with imperfect public monitoring originated by Abreu, Pearce, and Stacchetti (1990). Many papers apply this framework to characterize equilibria of repeated games that arise in applications. For example, Athey and Bagwell (2001), and Athey, Bagwell, and Sanchirico (2004) study the cartel problem, and Sannikov (2007) studies dynamic games in continuous time. Compared to these papers, the current paper focuses on the principal-agent relationship, in which the long-run limit is particularly interesting as it is related to the issue of persistent performance differences that has been actively studied in organizational economics recently.³

The problem of convergence has been studied in macroeconomics in the context of the long-run economic growth. Important contributions related to this paper include Green (1987), Spear and Srivastava (1987), Thomas and Worrall (1988 and 1990), and Kocherlakota (1996). Tools developed in these papers are used to prove stronger modes of convergence than convergence in distribution, such as convergence in probability and almost sure convergence. In the baseline model without limited liability considered in this paper, the limit distribution is non-degenerate, and the state does not converge in probability, so tools developed in the above papers are not likely to apply.

When the common discount factor of the principal and the agent is close to one, very general and precise characterization of the equilibrium payoff set can be given. This folk-theorem literature includes the classic papers by Fudenberg and Levine (1994) and Fudenberg, Levine and Maskin (1994). More recent installments include Escobar and Toikka (2013), Hörner, Takahashi, and Vieille (2013), and Barron (2013). However, these characterizations do not always offer useful insights for discount factors far below one, which are central in this paper.

Another related literature studies relational contracts, including Macleod and Malcolmson (1989), Baker, Gibbons and Murphy (1994, 2002), and Levin (2003), among others. This literature often assumes that both the principal and the agent are risk-neutral and have deep pockets, and predicts stationary optimal equilibria. In contrast, this paper explicitly assumes that the agent is risk-averse or is protected by limited liability, and explores the implications of these assumptions on equilibrium dynamics.

³ See Syverson (2011) and Gibbons and Henderson (2012) for surveys.

The optimal use of ex post public information about the agent's cost of effort has been studied in the mechanism design literature by Riordon and Sappington (1988), and Crémer and McLean (1988). Assuming transferable utilities, these papers show that the first-best outcome can be implemented even if signals about the environment or the correlations among players' types are far from perfect. Recent developments in this literature allow for non-transferable utilities. For example, Gary-Bobo and Spiegel (2006) studies the setting with limited liability. The current paper shows how insights from this literature help us understand the dynamics of repeated principal-agent interactions without commitment.

There is a recent literature of optimal dynamic contracts, although unlike this paper, most papers in this literature assume that the principal can commit to a long-term contract. For example, Sannikov (2008) predicts separation and retirement on the equilibrium path, but retirement can happen only when with commitment. DeMarzo and Sannikov (2006) and Biais, Mariotti, Plantin, and Rochet (2007) study optimal dynamic financial contracts. Zhu (2013) discusses a class of contracts (called the Quiet-Life contracts) in which the agent shirks from time to time on the equilibrium path. Similar to his result, shirking is also used as a reward in this paper, but here such shirking is always combined with cash rewards, while in Zhu's Quiet-Life contract the principal never rewards the agent through cash payment.

Various papers analyze the determinants of a player's continuation payoff other than his performance measure. For example, the reputation literature argues that the response to a player's performance also depends on the belief about his type. However, papers such as Ghosh and Ray (1996) assume that one of the two types is behavioral and the relationship is immediately terminated following the revelation of a bad type. It is not clear how these models can be generalized to accommodate changing environments. Similarly, Watson (1999) and Watson (2002) proposes a model which allows the bad type to cooperate to a certain degree, but these papers require that the scale of the cooperation is increasing over time before the relationship is terminated. There is little empirical evidence that a firm's size grows until the dismissal of its CEO though.

Another possible driving force of dynamics is the persistence in the exogenous environment. For example, McAdams (2011) studies future equilibrium path on a state. However, he assumes that both the state and players' actions are commonly observable

and conclude that the players' joint future equilibrium payoff is higher if the current state is better. This result does not explain the dependence of payment-performance sensitivity on the environment. The reputation literature

Two closely related papers, Fong and Li (2012) and Li and Matouschek (2013), also study equilibrium dynamics with non-transferrable utilities. Fong and Li (2012) analyzes dynamics in a principal-agent model with limited liability assuming that the agent is risk neutral and the output is binary. The current paper focuses on the case where the agent is risk averse, but also considers the risk neutral case, where it generalizes their result to a setting where the agent may be risk averse for low consumption. Li and Matouschek (2013) studies a principal-agent model in which the effort of the agent (the firm's manager in their paper) is observed. The observability of the agent's effort makes the long-run behavior of their model different from the long-run behavior of the model with moral hazard considered in this paper.

1.2 The Model

A principal (she) and an agent (he) interact repeatedly at time $t = 0, 1, 2, \dots$. In each period, the agent chooses an unobservable effort with a cost depending on an exogenous state variable Θ_t , which is privately observed by the agent. Specifically, the following events happen in period t :

1. The principal proposes a payment b_t to the agent. A negative b_t means a payment of $-b_t$ from the agent to the principal.
2. The agent accepts or rejects the payment.
3. The agent observes a state variable $\Theta_t \in \{B, G\}$, and chooses effort $e_t \in \{0, 1\}$.
4. The output X_t and a signal Z_t are publicly observed. X_t depends only on e_t , and Z_t depends only on Θ_t .

The stage game can be interpreted as a short-term "contract" with an unverifiable outcome. Under this interpretation, b_t is the wage payment specified in the contract. In the literature, it is usually assumed that when the agent rejects the contract, the stage game ends. This assumption is not important in the current setting, as in the optimal

equilibrium, the agent always chooses zero effort after rejecting the contract. In many applications, the agent is protected by limited liability, which adds the constraint that $b_t \geq 0$. The assumption that neither output X_t nor the signal Z_t is verifiable implies that any reward or punishment must be implemented through future interactions. Since the principal cannot commit to a long-run contract, there is an endogenous bound on the amount of reward and punishment that she can impose: if she promises too high a reward, she would prefer walking away to honoring the promise.

The above model will be referred to as the "baseline model" or the "model without limited liability". In many applications, the agent is protected by limited liability, so b_t cannot be negative. The model with this additional constraint is referred to as the "model with limited liability". Most of the analysis applies equally well to both models, and it is interesting to compare their dynamics, so these two models will be developed in parallel.

In what follows, the state variable Θ_t is often referred to as the "environment", and is assumed to be i.i.d. over time. Let μ_θ be the probability that $\Theta_t = \theta$ for $\theta \in \{B, G\}$. The i.i.d. assumption means that although there is fluctuation in the environment, it is not persistent. For simplicity, both the environment and the effort are assumed to be binary, where B and G mean bad and good environment respectively, and 0 and 1 mean "no effort" and "effort" respectively. This assumption will be relaxed in Section 6.

The environment does not affect the distribution of output X_t conditional on effort e_t ; it affects only the agent's cost of exerting effort. There is an *ex post* public signal Z_t about the environment that allows the principal to adapt incentives to the environment. It is assumed that both X_t and Z_t have probability densities with monotone likelihood ratios. Specifically, X_t has density $g(x, e)$ conditional on the effort e , and Z_t has density $f(z, \theta)$ conditional on the environment θ . Moreover, $g(x, 1)/g(x, 0)$ is non-decreasing in x and $f(z, G)/f(z, B)$ is non-decreasing in z . The assumption that the signal about the environment and the signal about the effort are separate makes the discussion of optimal response to the environment and the comparison to the empirical evidence easier. Section 6 discusses what happens when this assumption is relaxed.

One application where the principal does not observe the agent's cost of effort but has access to an *ex post* noisy measure of it is Macchiavello and Morjaria (2012)'s study of Kenyan flower exports during a period of domestic conflict. In the relationship between

a Kenyan flower exporter (the agent) and a foreign buyer (the principal), b_t is the price of flowers for the coming season t , and e_t is the exporter's effort of hiring and monitoring labors to pick and process flowers. The cost of labor fluctuates over time, and will be captured by the environment Θ_t , which is privately observed by the exporter and potentially different for different exporters. On the other hand, the foreign buyer may have access to some public indicator Z_t on the environment Θ_t , which may reflect whether there is a conflict in the country that causes exporters to lose their labor forces and whether the exporter is located in the region directly affected by the conflict.

Both the principal and the agent discount future payoff by $\delta \in (0, 1)$. The principal's average payoff is $(1 - \delta) \sum_t \delta^t (X_t - b_t)$ and the agent's is $(1 - \delta) \sum_t (u(b_t) - c_{\Theta_t} e_t)$. The function u represents the agent's utility from consuming the wage payment and is assumed to be a strictly increasing, twice differentiable and weakly concave function. The second term in the agent's payoff, $-c_{\Theta_t} e_t$, represents the cost of his effort. Effort level "zero" (or no effort) is costless, while the cost of effort level "one" depends on the environment. It is assumed that $c_B > c_G > 0$. Since the cost of effort is always positive, the stage game has a unique Bayesian Nash equilibrium in which no payment or effort happens. We normalize the two players' payoff functions so that both of them receive payoff zero in this equilibrium. In particular, $E[X_t | e_t = 0] = 0$, and $u(0) = 0$. Let $y = E[X_t | e_t = 1]$, the expected output when the agent exerts effort.

The solution concept is Perfect Public Equilibrium (PPE). Public randomization is allowed explicitly. Specifically, an i.i.d. random variable $\eta_t \sim U[0, 1]$ is publicly observed after payment b_t is made and before the agent chooses his effort.⁴ The public history at the beginning of period t , h^{t-1} , contains $\{(b_\tau, \eta_\tau, X_\tau, Z_\tau)\}_{\tau=0}^{t-1}$. The wage offer b_t and the agent's decision of acceptance or rejection depends on h^{t-1} , and the agent's effort depends on h^{t-1} , Θ_t , and η_t . Alternatively, one can imagine that the agent chooses a map from the environment to the effort based on h^{t-1} and η_t in Period t . This map will be denoted by $e_t = (e_{Bt}, e_{Gt})$, where $e_{\theta t}$ is the effort the agent plans to choose if $\Theta_t = \theta$. Since the signals about the principal's action and the agent's action are separate, the game has the product structure defined in Fudenberg and Levine (1994). As a result, every sequential equilibrium is payoff equivalent to a PPE.

⁴There is no need to introduce public randomization for the payment b_t as the principal always prefers paying $u^{-1}(E[u(\bar{b}_t)])$ to paying a random payment \bar{b}_t .

1.3 Preliminaries

This section provides some characterizations of the equilibrium frontier. After the discussion of the construction of the equilibrium frontier in Section 3.1, the main characterization results are proved in Section 3.2. There are two parallel but related ways to proceed from that section. Section 4 studies its implication to the short-run dynamics and characterizes the optimal informal incentives. It explains some empirical findings that are inconsistent with conventional wisdom. Section 5 studies the long-run dynamics, focusing on the possible types of long-run behaviors of the state variable w and how its long-run behavior depends on the agent's utility function u .

1.3.1 The equilibrium frontier

The equilibrium frontier $V(w)$ is defined as the principal's maximum payoff from an equilibrium that gives the agent payoff w . An equilibrium is called *optimal* if the principal's and the agent's expected payoffs are $(w, V(w))$ for some w . This paper focuses on the dynamics in optimal equilibria. Since the agent can always guarantee himself a payoff of zero by rejecting any negative payment b and never exerting any effort, there is no equilibrium in which the agent's payoff is negative. On the other hand, feasibility and individual rationality impose an upper bound on the agent's payoff. Let \bar{w} be the agent's maximum equilibrium payoff. Then V is defined on $[0, \bar{w}]$. It is easy to show that $V(\bar{w}) = 0$.⁵ Due to public randomization, V is concave on $[0, \bar{w}]$, and the equilibrium payoff set is $\{(w, v) : 0 \leq v \leq V(w)\}$.

Since the payment b_t is publicly observed, in analyzing the optimal equilibria, it can be assumed without loss of generality that all deviations regarding b_t cause the permanent switch to the static Bayesian Nash equilibrium, and the principal's continuation payoff should always be on the frontier without deviations regarding b_t . In other words, if the agent's expected continuation payoff at the beginning of Period t is w_t , then the principal's expected continuation payoff is $V(w_t)$ unless deviations regarding $b_{t'}$ happened for some $t' < t$. In what follows, a history is called *on the equilibrium path* if it contains no deviation regarding payment b .

⁵Suppose that $V(\bar{w}) > 0$. Then the principal can always give the agent equilibrium payoff $\bar{w} + \epsilon$ for sufficiently small $\epsilon > 0$ by increasing the up front payment b_0 in the first period.

To construct the program that characterizes V , it is helpful to start with the principal's maximum payoff $\mathcal{F}_e V$ when she induces the agent to exert a given effort profile $\mathbf{e} = (e_B, e_G)$:

$$\mathcal{F}_e V(w) = \max_{w_c, h} (1 - \delta)(-u^{-1}(h) + E[X|\mathbf{e}]) + \delta E[V(w_c(X, Z))|\mathbf{e}]; \quad (1.3.1)$$

$$s.t. \quad (1 - \delta)(h - E[c_{\Theta} e_{\Theta}]) + \delta E[w_c(X, Z)|\mathbf{e}] = w; \quad (1.3.2)$$

$$(2e_{\theta} - 1)\delta\{E[w_c(X, Z)|\theta, e = H] - E[w_c(X, Z)|\theta, e = L]\} \geq (1 - \delta)c_{\theta}, \text{ for } \theta \in \{B, G\}. \quad (1.3.3)$$

In the program, $w_c(x, z)$ is the agent's continuation payoff if the realization of (X, Z) is (x, z) , and $h = u(b)$ is the agent's utility from consuming the payment. It is more convenient to use h as the choice variable instead of b since doing this makes both constraints in the above program linear in choice variables. The first constraint, Eq. (1.3.2), is the principal's promise-keeping constraint. It states that the agent receives the promised expected payoff w if he chooses the effort profile \mathbf{e} . Eq. (1.3.3) is the agent's incentive-compatibility constraint, which states that the agent's gain in continuation payoff by exerting effort is no less than (no more than) the cost of effort in environment θ if he is supposed to exert effort (not to exert effort, respectively) in that environment.

The principal's maximum payoff is obtained by "concavifying" $\max_{\mathbf{e}} \mathcal{F}_e V(w)$. Specifically, the principal's maximum payoff $\mathcal{F}V(w)$ is given by

$$\mathcal{F}V(w) = \max_{e', e'', w' \leq w, w'' > w} \frac{w'' - w}{w'' - w'} \mathcal{F}_{e'} V(w') + \frac{w - w'}{w'' - w'} \mathcal{F}_{e''} V(w''). \quad (1.3.4)$$

Since the principal cannot commit not to walk away, her continuation payoff cannot be negative. Therefore, we set $\mathcal{F}V(w) = -\infty$ if the objective function on the right hand side of Eq. (1.3.4) is always negative. By Abreu, Pearce, and Stacchetti (1990), the equilibrium frontier V can be found by iterating \mathcal{F} starting from the feasible and individually rational payoff frontier. The frontier decreases over the iteration, and the limit is the equilibrium frontier V : $\mathcal{F}V = V$.

Figure 1 shows an equilibrium frontier generated by numerical simulation under the limited liability assumption. Figure 2 shows the $\mathcal{F}_e V$'s corresponding to this equilibrium frontier. By comparison, Figure 3 shows an equilibrium frontier generated by numerical simulation without limited liability, and Figure 4 shows the corresponding $\mathcal{F}_e V$'s.




FIGURE 1.1: The equilibrium frontier under limited liability

$u(b) = \sqrt{2b}$, $c_B = 3$, $c_G = 1$, $\mu_B = 0.2$, and $y = 15$. X_t and Z_t are both binary. $\text{Prob}(X_t = h|e_t = 1) = \text{Prob}(X_t = l|e_t = 0) = 0.8$, and $\text{Prob}(Z_t = g|\Theta_t = G) = \text{Prob}(Z_t = b|\Theta_t = B) = 0.7$.




FIGURE 1.2: Branches of the frontier under limited liability

There are several features worth noticing. First, the equilibrium frontier is monotonic in the example without limited liability (Figure 3), and is not monotonic in the example with limited liability (Figure 1). Under limited liability, $V(0) = 0$, and when w is close to zero, it is not possible to induce any effort from the agent. Intuitively, for w close to zero, the punishment that the principal can impose on the agent for low output is very limited, so incentives must be primarily provided through reward. This means that the agent receives rent because he can also receive the reward (although with smaller probability) when he does not exert effort. As w approaches zero, the rent eventually exceeds w , implying that it is not possible to induce effort. Indeed, repeating the static




FIGURE 1.3: The equilibrium frontier without limited liability
 $u(b) = 1 - \exp(-b)$, $c_B = 0.367$, $c_G = 0.2$, $\mu_B = 0.2$, and $y = 1$. X_t and Z_t are both binary. $\text{Prob}(X_t = h|e_t = 1) = \text{Prob}(X_t = l|e_t = 0) = 0.8$, and $\text{Prob}(Z_t = g|\Theta_t = G) = \text{Prob}(Z_t = b|\Theta_t = B) = 0.7$.

equilibrium is the unique PPE in which the agent receives expected payoff $w = 0$. This does not happen when the agent is not protected by limited liability since the principal can then first ask for a payment from the agent. Formal arguments will be given in Section 3.2.

Secondly, the optimal effort profile \mathbf{e} depends on the agent's expected payoff w (Figure 2 and 4). To provide incentives, the agent's expected continuation payoff w_t must change over time. Consequently, the equilibrium effort profile \mathbf{e} may also change over time. When there are two organizations with the same production function but independent draws of output conditional on the agents' efforts,⁶ the two agents' expected continuation payoff w_t will differ, which leads to different effort profile and thus different expected output between the two organizations. In other words, performance difference can result from noise in output without any persistent uncertainty about the production function. An interesting question is whether this type of performance different can persist in the

⁶Whether the environmental shock Θ_t is common to the two organizations is not important for this discussion




FIGURE 1.4: Branches of frontier without limited liability

long run or the two organizations' performances will converge over time. This is the central question in our study of the long-run behavior of the equilibrium in Section 5.

Thirdly, the curves $\mathcal{F}_e V$ intersect each other transversally in Figures 2 and 4. This implies that the transition from one effort profile to another is through a range in w on which V is linear. In general, the optimal equilibrium involves public randomization on a linear segment of V . The transversality of the intersections will play an important role in Section 5.

1.3.2 The first-order approach

To prepare for the study of short-run dynamics in Section 4 and long-run dynamics in Section 5, this subsection first proves the differentiability of the equilibrium frontier V , and then shows that the program Eqs. (1.3.1)-(1.3.3) can be solved using the first-order approach.

If repeating the static equilibrium is the only PPE of the repeated game, there will be nothing to characterize. This happens when the agent's incentive-compatibility constraint Eq. (1.3.3) can never be satisfied for $e_G = 1$. The following assumption will be maintained throughout to ensure that Eq. (1.3.3) can be satisfied with strict inequality at least in the good environment:

Assumption 1. $(1 - \delta)^{-1} \delta \bar{w} \int [g(x, 1) - g(x, 0)]^+ dx \in (c_G, c_B) \cup (c_B, \infty)$.

Here $\bar{w} = \sup\{w : V(w) \geq 0\}$ is the agent's maximum equilibrium payoff. The strongest incentive that the principal can provide is the left hand side of the inequality: the agent's continuation payoff w is zero if the realization of output, x , is such that $g(x, 1) \leq g(x, 0)$ and $w = \bar{w}$ if $g(x, 1) > g(x, 0)$. The first-order approach is valid when there exists some choice variables that satisfy Eq. (1.3.3) with strict inequality. For this reason, Assumption 1 rules out the possibility that the left hand side equals c_G or c_B .

In the program Eqs. (1.3.1)-(1.3.3), the choice variables are w_c and h . It is therefore unsurprising that some characterization results are stated in terms of w_c , the agent's continuation payoff. However, this quantity is usually not empirically observable. The following proposition relates the agent's expected continuation payoff w_t at the beginning of Period t to the payment b_t in that period, which is more likely to be observable by econometricians. A by-product is the differentiability of equilibrium frontier V .

Proposition 1. Assume that the agent is not protected by limited liability. Then V is differentiable on $(0, \bar{w})$, and in every optimal equilibrium, $V'(w_t) = -1/u'(b_t)$ on the equilibrium path if $w_t \in (0, \bar{w})$.

Proof. Consider the program Eq. (1.3.1)-(1.3.3) at some $w \in (0, \bar{w})$. The principal can give the agent expected payoff $w + \epsilon$ instead by adjusting h . Therefore,

$$V(w + \epsilon) \geq V(w) + (1 - \delta)u^{-1}(h(w)) - (1 - \delta)u^{-1}(h(w) + (1 - \delta)^{-1}\epsilon), \quad (1.3.5)$$

where $h(w)$ is the optimal choice of h in the program at w . The right hand side is differentiable in ϵ , and the derivative at $\epsilon = 0$ is $-1/u'(u^{-1}(h(w)))$. Now the concave function V is bounded from below by a differentiable function in a neighborhood of w , and the bound is tight at w . Therefore, V is differentiable at w with derivative $-1/u'(u^{-1}(h(w)))$. Notice that $u^{-1}(h(w))$ is the optimal payment at w . \square

The proof is a standard envelope-theorem argument. When u is strictly concave, $V'(w_t)$ is weakly decreasing in w_t , and $u'(b_t)$ is positive and strictly decreasing in b_t , so b_t is weakly increasing in w_t . This is intuitive: if the principal promises the agent higher continuation payoff, she starts paying it off by making a higher up-front payment. The

proposition also implies that $V' < 0$, so the equilibrium frontier is monotonically decreasing. In particular, $V(0) > 0$.

Now we are ready to write down the first-order condition for the program Eqs. (1.3.1)-(1.3.3).

Lemma 1. Assume that the agent is not protected by limited liability and that Assumption 1 holds. Consider a $w \in [0, \bar{w}]$ such that $V(w) = \mathcal{F}_e V(w)$ for some e at w . Then there exist κ_B and κ_G such that

$$V'(w_c(x, z)) - V'(w) = -\frac{\kappa_G f(z, G) + \kappa_B f(z, B)}{\mu_G g(x, e_G) f(z, G) + \mu_B g(x, e_B) f(z, B)} [g(x, 1) - g(x, 0)]. \quad (1.3.6)$$

Proof. In the program Eqs. (1.3.1)-(1.3.3), the objective function is concave in the choice variables and the constraints are linear in the choice variables. Therefore, the strong duality holds as long as there exists some w_c and h that satisfies Eq. (1.3.2) and satisfies Eq. (1.3.3) with strict inequality. (cf Proposition 5.3.1 of Bertsekas (1999).) By Assumption 1, there exists w_c that satisfies Eq. (1.3.3) with strict inequality. One can then choose h so that Eq. (1.3.2) is satisfied. Therefore, there exist Lagrange multipliers γ for Eq. (1.3.2) and $(2e_\theta - 1)\kappa_\theta$ for Eq. (1.3.3) so that

$$\begin{aligned} \frac{1}{w'(w^{-1}(h))} - \gamma &= 0; \\ V'(w_c(x, z)) + \gamma &= -\frac{\kappa_G f(z, G) + \kappa_B f(z, B)}{\mu_G g(x, e_G) f(z, G) + \mu_B g(x, e_B) f(z, B)} [g(x, 1) - g(x, 0)]. \end{aligned}$$

Proposition 1 and the first equation implies that $\gamma = -V'(w)$. \square

When $V(w) > \mathcal{F}_e V(w)$ for all e , a public randomization is strictly optimal. The solution to the program Eqs. (1.3.1)-(1.3.3) does not characterize the true evolution of the agent's continuation payoff. Therefore, the above lemma only considers the case when a particular e is optimal and there is no need of public randomization. In Eq. (1.3.6), κ_B and κ_G are the Lagrange multipliers of the constraint Eq. (1.3.3) for $\theta = B$ and $\theta = G$, respectively. The Lagrange multiplier of the promise-keeping constraint Eq. (1.3.2) proves to be $-V'(w)$. Eq. (1.3.6) should be understood to allow for corner solution in the following sense: $w_c(x, z) = \bar{w}$ if and only if the right hand side of Eq. (1.3.6) is less than or equal to $V'(\bar{w}) - V'(w)$, and $w_c(x, z) = 0$ if and only if the right hand side of Eq. (1.3.6) is greater than or equal to $V'(0) - V'(w)$.

Now consider the case where the agent is protected by limited liability. As one can see from Figure 1, the equilibrium frontier is not monotonic when the agent is protected by limited liability. In general, the condition that $V'(w_t) = -1/u'(b_t)$ does always hold. The following assumption will be helpful in establishing the differentiability of V :

Assumption 2. $g(x, 1)/g(x, 0)$ and $f(z, G)/f(z, B)$ are strictly increasing and continuous in x and z , respectively.

Proposition 2. Assume that the agent is protected by limited liability. Under Assumptions 1 and 2, the following holds:

1. V is differentiable on $(0, \bar{w})$. If on the path of an optimal equilibrium, $V'(w_t) < -1/u'(0)$, then $V'(w_t) = -1/u'(b_t)$; if $V'(w_t) \geq -1/u'(0)$, then $b_t = 0$.
2. Eq. (1.3.6) still holds when $V(w) = \mathcal{F}_e V(w)$ for some e .
3. $V(0) = 0$ and V is linear in a neighborhood of zero, and the unique PPE with $w = 0$ is repeating the static equilibrium forever.

Proof. See Appendix A. □

In fact, the differentiability of V when the right derivative of V is smaller than $-1/u'(0)$ can be derived as in Proposition 1. Assumption 2 helps us prove the differentiability of V at smaller w . In that case, the limited liability constraint is binding, but the principal can adjust reward for high output to accommodate change in the agent's expected payoff. The numerical simulation shown in Figure 1 does not satisfy Assumption 2, as both the output X and the signal about the environment Z are binary, and the equilibrium frontier there is not differentiable. Assumption 2 will always be maintained when we study the limited liability case.

Notice that when the agent is protected by limited liability, $b_0 = 0$ if $V'(w_0) = 0$. In other words, there is no up-front payment in the principal-optimal equilibrium. Intuitively, the principal has incentive to delay cash payment when the agent is protected by limited liability. The final assertion of the proposition opens up the possibility of "termination": when w_t reaches zero, the principal and the agent will repeat the static equilibrium from then on, and in this case we say that their relationship is "terminated". Of course, we do not know yet whether w_t will reach zero with positive probability on the equilibrium

path; this will be addressed in Section 5. Recall that when the agent is not protected by limited liability, $V(0) > 0$, and the payoff pair never leaves the frontier on the path of an optimal equilibrium, so termination does not happen in any optimal equilibria without limited liability.

Starting from Eq. (1.3.6), one can study how w_c depends on x and z , and how the agent's effort in the next period depends on w_c in turn. The key step is to find out the signs of κ_B and κ_G . This will be done in the next section. Alternatively, one can treat Eq. (1.3.6) as the difference equation that describes the transition from w_t to w_{t+1} , and then study the long-run behavior of the stochastic dynamical system. This approach pays more attention to the probability distribution of $w_c(X, Z)$ than to the function w_c itself, and will be taken up in Section 5.

1.4 Optimal informal incentives

1.4.1 Optimal response to the fluctuating environment

This subsection studies how $w_c(x, z)$ depends on x and z and compares the result with some empirical evidence. To gain intuition consider first the risk-neutral benchmark where $u(b) = b$.

By Proposition 1, $V'(w) = -1$ for all $w \in [0, \bar{w}]$. Using Eq. (1.3.2), we can rewrite the objective function in the program Eqs. (1.3.1)-(1.3.3) as

$$(1 - \delta)E[X - c_\Theta e_\Theta | e] + \delta V(0) - \delta w.$$

None of the choice variables appear in this new objective function. As a result, the optimal choice of $w_c(x, z)$ can be made independent of z :⁷ when $e_B = e_G = 1$, choose $w_c(x)$ so that

$$\delta \int w_c(x)[g(x, 1) - g(x, 0)]dx = (1 - \delta)c_B;$$

when $e_B = 0$ and $e_G = 1$, choose $w_c(x)$ so that

$$\delta \int w_c(x)[g(x, 1) - g(x, 0)]dx = (1 - \delta)c_G.$$

⁷The case $e_B = e_G = 0$ is trivial. The choice of $e_B = 1$ and $e_G = 0$ can never be optimal: it is dominated by $e_B = e_G = 1$ if $y > c_G$, and it is dominated by $e_B = e_G = 0$ if $y < c_B$.

In words, when the principal wants the agent to exert effort in both environments, the agent gains $(1 - \delta)c_B$ in future payoff (and loses $(1 - \delta)c_\theta$ immediately) by exerting effort in either environment; when the principal wants the agent to exert effort only in the good environment, the agent gains $(1 - \delta)c_G$ in future payoff by exerting effort in either environment. In the former case, the agent receives "rent" in the good environment, as he strictly prefers exerting effort to not exerting effort, and in the latter case, the principal creates an unnecessary incentive of $(1 - \delta)c_G$ in the bad environment. The risk-neutral case is special in that neither of these is costly.

When the agent is risk-averse, V will not be linear. Then the principal wants to reduce the spread of the agent's continuation payoff as much as possible while respecting his incentive-compatibility constraint. When $e_B = e_G = 1$, this means that the principal wants to reduce the strength of incentive in the good environment compared to the risk-neutral benchmark. When $e_B = 0$ and $e_G = 1$, reducing the spread in w_c calls for minimizing the agent's "incentive" in the bad environment, or "insuring" him against noise in output in the bad environment.

To reduce the agent's incentive in Environment θ from the risk-neutral benchmark, the principal can reduce $w_c(x, z)[g(x, 1) - g(x, 0)]$ for some z . However, this also reduces the agent's incentive in the other environment. To respect the agent's incentive-compatibility constraint in the other environment, the principal has to raise $w_c(x, z)[g(x, 1) - g(x, 0)]$ for some other z . Intuitively, the optimal adjustment reduces $w_c(x, z)[g(x, 1) - g(x, 0)]$ for those z that suggest that the environment is likely to be θ , while raises $w_c(x, z)[g(x, 1) - g(x, 0)]$ for other z . This intuition is confirmed by the following proposition, which is derived from the first-order condition Eq. (1.3.6). In stating this proposition, it is convenient to call an output level x "high" if $g(x, 1) > g(x, 0)$ and "low" if $g(x, 1) < g(x, 0)$.

Proposition 3. Under Assumptions 1 and 2, for every $w_0 \in [0, \bar{w}]$, there exists an optimal equilibrium that gives the agent expected payoff w_0 with the following property:

- On the equilibrium path, $w_{t+1} \geq w_t$ if output X_t is "high", and $w_{t+1} \leq w_t$ if X_t is "low", as long as $e_{Gt} = 1$;
- If $e_{Bt} = e_{Gt} = 0$ at this history, then w_{t+1} is independent of X_t and Z_t .

- If $e_{Bt} = 0$ and $e_{Gt} = 1$ at this history, then w_{t+1} is weakly increasing in Z_t if X_t is high, and w_{t+1} is weakly decreasing in Z_t if X_t is low;
- If $e_{Bt} = e_{Gt} = 1$ at this history, then w_{t+1} is weakly decreasing in Z_t if X_t is high, and w_{t+1} is weakly increasing in Z_t if X_t is low;

In any optimal equilibrium, w_{t+1} is independent of Z_t for all realizations of X_t only if V is linear on the convex hull of $\{w_{t+1}(x_t, z_t)\} \cup \{w_t\}$

An output x is high when the likelihood ratio $g(x, 1)/g(x, 0)$ is greater than one. The first assertion of the proposition is that the principal rewards the agent by raising his expected continuation payoff when he produces a high output, and punishes him when he produces a low output.⁸ In fact, when $e_{Gt} = 1$, both κ_G and κ_B in Eq. (1.3.6) are non-negative, so $V'(w_c(x, z)) - V'(w)$ and $g(x, 1) - g(x, 0)$ have the opposite signs. Since V' is decreasing, this implies that $w_c(x, z) - w$ and $g(x, 1) - g(x, 0)$ have the same sign.

When the agent is not expected to exert effort in any environment, the principal does not want any spread in his continuation payoff, so w_{t+1} should be constant. When the agent is expected to exert effort only in the good environment, the principal reduces his incentive in the bad environment compared to the risk-neutral benchmark by offering the agent higher reward for high output and harsher punishment for low output when the signal Z_t is high, which suggests that the environment is likely to be good. In fact, in this case the agent's incentive-compatibility constraint is not binding in the bad environment, so $\kappa_G \geq 0$ and $\kappa_B = 0$ in Eq. (1.3.6). This leads to the desired monotonicity in Z_t . When the agent is expected to exert effort in both environments, the same intuition suggests the opposite monotonicity in Z_t . In this case, $\kappa_B \geq \mu_B \kappa_G / \mu_G$, and the monotonicity of w_{t+1} in Z_t follows.

As mentioned in Section 3.2, the agent's continuation payoff is often not observable by the econometrician. Fortunately, when u is strictly concave, Propositions 1 and 2 allow us to translate the monotonicity of w_{t+1} in Z_t directly to the monotonicity of b_{t+1} in Z_t . The comparison between w_{t+1} and w_t can be similarly translated to the comparison between b_{t+1} and b_t . Therefore, if the econometrician groups observations by the level of output, and regresses bonus payment at the beginning of Period $(t + 1)$ and at the end

⁸This result may not be true when $e_{Bt} = 1$ and $e_{Gt} = 0$, since in this case κ_G in Eq. (1.3.6) may be negative.

of Period t on a public signal about the environment at Period t , he should coefficients of the opposite signs for high and low outputs. Moreover, the monotonicity indicates whether the agent is expected to exert effort in the bad environment.

The problem with the above regression scheme is that since output is a continuous variable, it is hard to group observations according to output. A more common design is to regress bonus payment on the output, the public signal about the environment, and their interaction. When the agent is expected to exert effort in both environments, Eq. (1.3.6) implies that $V'(w_c(x, z))$ has increasing differences in x and z . When $u(b) = \log b$, Proposition 1 implies that $V'(w) = -b$, so the principal's payment at the end of Period t and at the beginning of Period $(t+1)$, b_{t+1} , as decreasing differences in x_t and z_t , and one expects a negative coefficient on the interaction term in the above difference-in-difference regression.

An example is the empirical study by Macchiavello and Morjaria (2012) of Kenyan flower exporters. It seems that the foreign buyers expected the exporters to fulfill their delivery commitment even when there was shock to cost of labor. They show that reliability during the violence positively correlate with their survival rate and volume in the future for exporters in the region that is affected by the violence (the "conflict region"), and such correlation is not significant for exporters in the region that is not affected by the violence (the "non-conflict region"). This is consistent with the theoretical prediction with high effort expected in both states. When the foreign buyer receives the signal that suggests a shock (which means that he learns that the exporter is in the conflict region), he tends to reward those exporters who deliver during the crisis with higher volume in the future, and punish those who do not deliver with lower volume and even termination of the relationship. When the signal suggests that there is no shock (which means that the exporter is in the non-conflict region), both the reward and the punishment are smaller.

Conventional wisdom holds that the risk-neutral principal should insure the risk-averse agent against exogenous shocks that can affect his output, so the punishment should be less harsh when ex post public information suggests that the environment is tough. However, the empirical findings of Macchiavello and Morjaria (2012) does not support this hypothesis.

1.4.2 Evolution of effort level

The previous subsection characterizes how the agent's continuation payoff evolves over time, as the optimal response to his output and information about the environment. To understand the evolution of his effort level, it remains to characterize its dependence on his continuation payoff. In general, this characterization is difficult to obtain, as the cost of giving Player 1 incentive not only depends on the slope of the equilibrium frontier, but also on how "concave" the frontier is in different regions. Indeed, Figure 2 shows that the dependence of the agent's effort on his continuation payoff may not be monotonic.

When both parties are sufficiently patient, the equilibrium frontier V and the $\mathcal{F}_e V$'s are close to the first-best limit. The dependence of the agent's effort e on his expected payoff w in the optimal equilibrium can be deduced from its counterpart in the first best. In the first best, the probability that the agent exerts effort in Environment θ , q_θ , solves the following program:

$$\begin{aligned} \max_{q_\theta, h} \quad & \sum_{\theta} \mu_{\theta} q_{\theta} y - u^{-1}(h) \\ \text{s.t.} \quad & h - \sum_{\theta} \mu_{\theta} q_{\theta} c_{\theta} = w. \end{aligned}$$

Since no incentive in continuation payoff is needed in the first best, the above program is a purely static. The objective function is the principal's expected payoff, and the constraint is the promise-keeping constraint, the analogue of Eq. (1.3.2). One can eliminate h and show that q_θ maximizes the following expression:

$$\sum_{\theta} \mu_{\theta} q_{\theta} y - u^{-1} \left(w + \sum_{\theta} \mu_{\theta} q_{\theta} c_{\theta} \right).$$

It is easy to see that $q_B > 0$ only if $q_G = 1$, and both q_B and q_G are decreasing in w . Therefore, as w increases, the optimal expected efforts (q_B, q_G) change from $(1, 1)$ to $(0, 1)$ and then to $(0, 0)$ under the first best, while one or two of these regimes may be missing. The transition from one regime to another is through a non-empty range of w for which a randomization is strictly optimal: $q_B \in (0, 1)$ or $q_G \in (0, 1)$.

By the Nash-threat folk theorem of Fudenberg, Levine, and Maskin (1994), the equilibrium frontier converges uniformly to the first-best frontier on $[w_l, w_h]$ for all $w_l > 0$ and $w_h < \bar{w}_{FB}$, where \bar{w}_{FB} is the highest payoff that the agent can receive in the first

best when the principal's payoff is non-negative. Therefore, on each such closed interval $[w_l, w_h]$, the effort (e_B, e_G) in the equilibrium changes from $(1, 1)$ to $(0, 1)$ and then to $(0, 0)$ as w increases in the optimal equilibrium, while one or two of these regimes may be missing, and the transition between two regimes is through a non-trivial range of w for which public randomization is strictly optimal in the first period.

The above discussion also makes it clear that this pattern of decreasing effort in continuation payoff comes from the fact that the agent has decreasing marginal utility of consumption. As the agent's expected payoff increases, his value of compensation decreases, and therefore the optimal effort decreases. Proposition 3 shows that the agent's continuation payoff increases when his output is high and decreases when his output is low. The combination of the above two results shows that an agent who has produced a series of high outputs will enter a region with high compensation but weak incentives. His expected output in the next period is thus lower. Conversely, an agent who has produced a series of low outputs will have strong incentives, and his expected output in the next periods will be higher. This mean-reversion behavior helps us understand why the effect of the initial state vanishes over time when we discuss the long run limit of the equilibrium path in the next section. Interestingly, effort may be decreasing in expected payoff in the optimal equilibrium even when a single effort level prevails in the first best for all w . In the numerical simulation shown in Figures 3 and 4, the parameters are chosen so that $(1, 1)$ is always optimal in the first best, while the agent's effort in the bad environment is still decreasing in w . Intuitively, the agent's marginal utility of consumption is always diminishing and it is always harder to provide incentives to an agent with high expected utility, regardless what the optimal effort level is in the first best.

In the context of corporate governance, decreasing effort in expected payoff means that successful CEOs enjoy temporary "easy life" with weak incentives yet high compensation. This finding is broadly consistent with the empirical literature of endogenous governance structure of firms. For example, Schoar and Washington (2011) shows that managers of the firms whose performance beat analysts' consensus forecast for multiple times are more likely to call for special shareholder meetings and propose changes in governance structure that are good for themselves and bad for share holders. A bad governance structure typically allows a manager more freedom to pursue his own interest and weakens his incentive to maximize the shareholders' value.

The mean-reversion discussed above breaks down in the model with limited liability. In this case, it is not possible to provide the agent with incentives when his continuation payoff is close to zero, so on the whole domain $[0, \bar{w}]$ of the frontier, the agent's effort always increases in w for w close to zero. As illustrated in Section 5.3, difference in the initial state may be amplified and lead to permanent difference in the long run in the model with limited liability.

Although it is hard to characterize the dependence of effort on the agent's expected payoff w for discount factor far below one, it is easy to characterize the range of w in which the agent does not exert effort in either environment.

Proposition 4. Suppose that $\mathbf{e} = (0, 0)$ is optimal for some $w \in [0, \bar{w}]$. If $V'(w) < 0$, then $\mathbf{e} = (0, 0)$ is optimal for all $\bar{w} \in [w, \bar{w}]$. If $V'(w) \geq 0$ (which can happen only in the model with limited liability), then $\mathbf{e} = (0, 0)$ is optimal for all $\bar{w} \in [0, w]$.

To understand the intuition behind this proposition, first consider the model without limited liability. Then the proposition says that the effort profile $(0, 0)$ can prevail in the optimal equilibrium only for w sufficiently close to \bar{w} . This result holds for all discount factor δ . Zero effort in both environments is the extreme example of a successful agent's "easy life" discussed above. The connection between diminishing marginal utility of consumption and increasing cost of incentive provision is easier to convert to a formal proposition in this extreme case. Notice that the agent's continuation payoff will be strictly lower than his current expected payoff when $(0, 0)$ is optimal, so eventually his expected payoff w leaves the region where zero effort prevails. In other words, the extreme easy life with zero effort is temporary. This is different from the permanent "retirement" in Sannikov (2008). In the model with limited liability, zero effort also prevails for w close to zero, as it is not possible to provide the agent with any incentives in this case, as discussed before.

Zhu (2013) studies the optimal dynamic contract in a continuous-time principal-agent model. He shows that the contract may allow the agent to shirk on the equilibrium path as a way to reward him, and calls such a contract a "quiet-life" contract. There is one important difference between his quiet-life contract and shirking in the optimal equilibrium characterized in Proposition 4. In a quiet-life contract, the principal's payment to the agent is constant over time and independent of whether the agent is supposed to exert effort. In the model considered in this paper, an "easy life" is a combination of

high compensation and weak incentive: the agent's compensation will be lower after his easy life ends.

1.5 Long-run Dynamics

In this section, we study the long-run behavior of optimal equilibria. In particular, we want to know whether two optimal equilibria with different payoff pairs $(w', V(w'))$ and $(w'', V(w''))$ will become close to each other in the long run, and if so, in what sense. The main result in Section 5.1 shows that under some assumptions, the agent's expected continuation payoff w_t converges to a unique distribution which is independent of its initial value w_0 . The result will not be very interesting if w_t cannot summarize the future path of the optimal equilibrium, so this issue will be addressed before the main result is stated. Section 5.2 sketches the proof of the main result. Besides presenting the result on the baseline model, the paper also studies how changes in assumptions affect the limit distribution of w_t . Specifically, Section 5.1 compares the limit distributions in models with and without limited liability, Section 5.3 studies the case where the agent is protected by limited liability but otherwise risk-neutral, and Appendix B studies the case where the principal can commit to a long-term contract. The differences between the limit distributions of different settings are discussed in the context of persistent performance difference: whether the performances of two organizations with independent shocks will differ in the long run. The main result will be generalized to a setting that allows for continuous effort and more than two environments in Section 6.

1.5.1 The main result

The main result on the long-run dynamics is concerned with the distribution of w_t , the agent's expected continuation payoff at the beginning of Period t . For convenience, we will refer to it as the "state" of the equilibrium. This should not be confused with the state of the environment Θ_t . An event that happens in the first period of an optimal equilibrium with the agent's expected payoff being w will be called an event "at w ". However, if two very different equilibria deliver the same payoff pair $(w_t, V(w_t))$, the "state" does not completely determine the future path of the equilibrium and is not

very useful. Therefore, it is desirable to have that the optimal equilibrium with payoff pair $(w, V(w))$ is unique for every $w \in [0, \bar{w}]$.

Unfortunately, this uniqueness result is not true in general. For example, we have seen that when the agent is protected by limited liability, $V(0) = 0$, and no effort can be induced for w close to zero. Let w_s be the lowest w such that no public randomization is needed and the agent exerts effort at least in one environment at w . Then for $w \in [0, w_s]$, one optimal equilibrium involves a public randomization between termination of the relationship ($w = 0$) and resuming the production at w . Alternatively, if $\delta^{-1}w \leq w_s$, playing the static equilibrium in the first period and starting the next period at $\delta^{-1}w$ can also be optimal, since this gives the agent and the principal payoff $(w, \delta V(\delta^{-1}w))$, and V is linear on $[0, w_s]$. Another possibility is that at an intersection $\mathcal{F}_{\mathbf{e}_1}V$ and $\mathcal{F}_{\mathbf{e}_2}V$ for some $\mathbf{e}_1 \neq \mathbf{e}_2$, the two effort profiles give the same payoff. However, unless the two curves are also tangent to each other at the intersection, a public randomization is strictly optimal at the intersection.

Let $W_I = \{w : V \text{ is linear in a neighborhood of } w\}$. In other words, $\{(w, V(w)) : w \in W_I\}$ is the set of all non-extreme points of the equilibrium frontier $\{(w, V(w)) : w \in [0, \bar{w}]\}$. Moreover, no public randomization is possible at $w \notin W_I$. We have seen that the optimal equilibrium at w may not be unique when $w \in W_I$. The next proposition shows that once w_t leaves W_I , it never returns to it.

Assumption 3. \bar{w} is not in the closure of W_I . When the agent is not protected by limited liability, 0 is not in the closure of W_I either.

Proposition 5. Suppose Assumptions 1-3 holds, $w \notin W_I$ and an effort profile \mathbf{e} is optimal at w . Then the optimal choice w_c in the program Eqs. (1.3.1)-(1.3.3) is such that $\dot{w}_c(x, z) \in W_I$ for (x, z) with Lebesgue measure zero. Furthermore, if both w_c and \tilde{w}_c are optimal in the program, then $w_c(x, z) = \tilde{w}_c(x, z)$ almost everywhere.

Proof. See Appendix A. □

When the continuation payoff pair in an optimal equilibrium is almost always an extreme point of the equilibrium payoff set, the model is said to have the "bang-bang" property. Abreu, Pearce and Stacchetti (1990) establishes a sufficient condition for the bang-bang property. For w not in the closure of W_I , κ_B and κ_G in Eq. (1.3.6) cannot both be zero, and the proof of the bang-bang property is essentially the same as APS's proof. The only

difference is that we do not require the likelihood ratio functions to be analytic; they only need to be continuous and strictly monotonic. The case where w is at the boundary of W_I is more complicated. In this case W_I has a connected component (w', w'') and $w = w'$ or $w = w''$. It is possible that $\kappa_B = \kappa_G = 0$, but since w is at the boundary of $[w', w'']$, it can be shown that $w_c(x, z)$ must be at the boundary $\{w', w''\}$ too for almost all (x, z) and is unique. The reason is that otherwise $\mathcal{F}_e V$ can be linearly extended beyond w , violating the assumption that w is at the boundary of W_I . This extension argument fails when $w = 0$ or \bar{w} . Therefore, we make Assumption 3. When the agent is protected by limited liability, 0 is in the closure of W_I , but we have seen that the unique equilibrium at zero is repeating the static equilibrium forever.

Assumption 3 is an assumption on an endogenous object, the equilibrium frontier V . Nevertheless, considering that there are standard algorithms for computing the equilibrium frontier (see Phelan and Townsend (1991)), being able to predict the long-run behavior of the state based on information about V is still valuable. In practice, it means that no public randomization is needed at both ends of the equilibrium frontier, and it is optimal to induce effort at least in the good environment when $w = \bar{w}$. When Assumption 3 fails, the main result, Theorem 1, still holds when we assume that whenever w is in a linear segment of V , the principal and the agent publicly randomize between the two end points of the linear segment.

By this proposition, as long as $w_0 \notin W_I$, $w_t \in W_I$ with probability zero for all t . This result allows us to ignore the possibility of $w_0 \in W_I$ for now and return to it after we establish the main result for all $w_0 \notin W_I$. Now the only possible multiplicity of optimal equilibrium comes from the multiplicity of the optimal effort profile e . We assume that it does not happen:

Assumption 4. For each $w \notin W_I$, the optimal effort profile e is unique at w , and will be denoted by $e(w)$.

As mentioned above, two different effort profiles e and e' are both optimal at w only if the curves $\mathcal{F}_e V$ and $\mathcal{F}_{e'} V$ intersect and are tangent to each other at w . Since this tangency is constrained to be true in the construction of V , and there are only finitely many curves to consider, it is conjectured that Assumption 4 holds for generic parameters. This conjecture remains to be checked.

Under Assumptions 1-4, the optimal equilibrium starting from each $w \in [0, \bar{w}] \setminus W_I$ is unique up to events with probability zero. Therefore, the equilibrium dynamics is characterized by the evolution of the state w_t . The final assumption needed for the main result is the following:

Assumption 5. u is strictly concave, and $g(x, 1)/g(x, 0)$ is bounded away from both zero and infinity.

It contains two parts. First, it assumes that the agent is strictly risk averse. If the agent is risk neutral over a wide range of payment level, then it may be possible to implement an optimal stationary relational contract such that the payment is always in that risk neutral range on the equilibrium path, and there is no dynamics. Section 5.3 discusses the case where the agent is protected by limited liability and otherwise risk neutral, and shows that the conclusion of Theorem 1 is not true there. Secondly, Assumption 5 requires that the output cannot be an arbitrarily precise signal about the agent's effort. This rules out incentive schemes similar to Mirrlees (1975). It will imply that no matter what the current state of the relationship is, the principal has to use both "carrot" and "stick" substantially to provide incentives in every optimal equilibrium.

Theorem 1. Under Assumptions 1 to 5, there exists a unique probability distribution π on $[0, \bar{w}]$ such that in every optimal equilibrium, the probability distribution of w_t converges to π in the following sense: $\sup_{A \in \mathcal{B}([0, \bar{w}])} |\text{Prob}(w_t \in A) - \pi(A)| \rightarrow 0$ as $t \rightarrow \infty$, where the supremum is taken over all Borel subsets of $[0, \bar{w}]$.

Roughly speaking, the theorem says that the difference in the initial condition w_0 vanishes in expectation as time approaches infinity. To build more intuition, consider two organizations with the same production function but independent draws of output conditional on the agents' efforts. The environmental shock Θ_t may or may not be common to both organizations. Since their draws of outputs are independent, even if they start with the same initial state w_0 , which may be determined by the market condition and the bargaining power of principals and agents, the state w_t will be different in the two organizations after a while. Since the agent's optimal effort level e depends on w as can be seen from Figure 2, the two agents may exert different efforts at each Period t , leading to different expected output. In other words, conditional on public information up to time t , the two organizations may have different expected output. This can be called a *transient performance difference*.

It is natural to ask whether the transient performance difference will persist over time. In other words, whether the organization with higher performance will continue to outperform the other organization, or the performances of the two organizations will converge over time. The theorem gives an answer under its assumptions: the expected output in Environment θ , $yE[e_\theta(w_t)]$, will converge to $y \int \pi(w)d\pi(w)$, as time t approaches infinity. In fact, the expectation of every bounded function a of the state w , $E[a(w_t)]$, will converge to $\int a(w)d\pi(w)$, a limit that is independent of the initial state. Moreover, since the dynamical system can be released at any time, the expectation of $a(w_t)$ conditional on public history up to time s also converges as $t \rightarrow \infty$ and s is kept fixed: $E[a(w_t)|h^s] \rightarrow \int a(w)d\pi(w)$. The probability distribution π will be referred to as the unique *invariant* distribution.

However, this theorem only asserts convergence of the distribution, and it does not imply convergence in probability unless the invariant distribution π is a point mass. In the two-organization example, convergence in probability means that the difference between their states $w_t^{(1)}$ and $w_t^{(2)}$ is close with high probability: $\text{Prob}(|w_t^{(1)} - w_t^{(2)}| > \epsilon) \rightarrow 0$ as $t \rightarrow \infty$. In fact, this probability will be bounded away from zero for every $\epsilon > 0$ in the baseline model where the agents are not protected by limited liability. In this case, it is hard to fully characterize the invariant probability distribution π , but the following proposition shows that π has a "full" support.

Proposition 6. The support of π is $[0, \bar{w}] \setminus W_I$, $\pi(\{0\}) > 0$, and $\pi(\{\bar{w}\}) > 0$.

The proof of this proposition relies on the lemmas that lead to Theorem 1, and thus will be presented after the proof of Theorem 1 in Appendix A. By this proposition, the limit distribution π is non-degenerate. If the shocks Θ_t and $X_t - E[X_t|e_t]$ in two organizations are independent, then the joint distribution of their states $w_t^{(1)}$ and $w_t^{(2)}$ are also independent, and thus in the long-run their states do not converge. Let $y_t^{(i)} = yE[e_t^{(i)}]$ be the expected output of Organization i at period t . Then unless the agent's effort level is constant over the whole state space $[0, \bar{w}]$,⁹ the two organizations' y_t will be different in the long-run. The fact that their states converge to a limit distribution π that is independent of the initial state implies that the expectation of their performance difference $y_t^{(1)} - y_t^{(2)}$ has expectation zero, but the variance of it is positive and converges to a constant. Therefore, there is performance difference in the long run, but the leader

⁹When the discount factor δ is close to one, the optimal effort level at each w is close to the optimal effort level at the feasible frontier at w .

will change over time so that no organization has a higher long-run expectation of y_t . Section 4.2 provides a heuristic discussion of the evolution of the effort level and why it may exhibit mean-reversion.

Figures 5 and 6 show simulation result on the probability distribution of w_t for the numerical example without limited liability. The equilibrium frontier has been shown in Figure 3. The parameters are chosen so that it is always optimal for the agent to exert effort in both environments in the first best. Figure 4 shows that the agent always exerts effort in the good environment in every optimal equilibrium. Therefore, the main concern is whether the agent exerts effort in the bad environment. When he does, the output of the organization is robust against shocks to the cost of effort.

Figure 5 shows the evolution of the expected effort in the bad environment over time for the two extreme initial states, w_0 and \bar{w} , which correspond to the principal-optimal equilibrium and the agent-optimal equilibrium respectively. For a discount factor of 0.7, convergence to the long-run limit is relatively quick, and the effect of the initial condition vanishes: two organizations that start from the principal's best state and her worst state have the same expected output in the long run. Figure 6 shows the cumulative distribution function of the invariant distribution π . As expected, it is independent of the initial state, has full support, and has point masses at w_0 and \bar{w} . Although both output X_t and the signal about the environment Z_t are binary, so the bang-bang property of Proposition 5 does not hold, the CDF is almost flat in the linear segment of V , and in this example the unique optimal equilibrium on the linear segment of V involves public randomization in the first period. Figure 6 illustrates that the convergence of expected effort in the bad environment to a long-run limit is not due to the convergence of w_t to a stationary state, and the process cannot enter a cycle either. The full-support result is not a consequence of binary effort choice, as Section 6 shows that the same result holds in a setting with continuous effort choice.

When the agent is protected by limited liability, the invariant probability distribution is clearly the unit mass at zero. This is because when $w_0 = 0$, $w_t = 0$ for all t , so the unit mass at zero is an invariant distribution, and the theorem says that it is the unique invariant distribution. If repeating the static equilibrium is interpreted as the termination of the relationship, the theorem implies that in a principal-agent relationship with limited liability, the relationship is eventually terminated with probability one. Notice




FIGURE 1.5: Evolution of the agent's effort in the bad environment

The agent always exerts effort in the good environment. The two curves correspond to two different initial states: $w_0 = 0$ and $w_0 = \bar{w} = 0.378$. The sample size is 10,000. The discount factor is $\delta = 0.7$.

that this result holds for all $\delta < 1$. Therefore, although $V(w)$ converges to the feasible and individually rational frontier as $\delta \rightarrow 1$ for $w > 0$, in the long run the state w will not stay close to w_0 even when δ is close to one. This suggests that the characterization of the long-run dynamics is not trivial even if one can characterize the equilibrium frontier. The folk theorem gives a good approximation of the equilibrium frontier for δ close to one, but it does not contain information about the long-run dynamics in this case.

Appendix B studies the long-run dynamics under the optimal long-term contract assuming that the principal can commit to such a contract, and shows that the limit distribution of state is very different from that of the no-commitment model discussed above. Under the assumption that the agent is not protected by limited liability and his marginal utility of consumption converges to zero as consumption goes to infinity,¹⁰ w_t eventually enters a "retirement" regime that was characterized in Sannikov (2008). In this regime, the principal makes a constant payment to the agent in each period, and the agent does not exert effort and enjoys a constant continuation payoff. The full support result or Proposition 6 and the long-run performance difference are gone. Notice that the principal's continuation payoff in the "retirement" regime is always negative, so this

¹⁰This assumption is satisfied by both the exponential utility and the power function utility.




FIGURE 1.6: Simulated cumulative distribution function of the invariant distribution

For each curve, all samples are released with the same initial value ($w_0 = 0$ for the blue curve and $w_0 = \bar{w}$ for the red curve), and the empirical CDF is computed at $t = 200$. The sample size is 10,000.

regime does not exist when the principal cannot commit to a long-term contract. Therefore, the long-run dynamics of a model with commitment is not a good approximation to its counterpart without commitment. If one is interested in the long-run behavior of a principal-agent relationship, he/she has to first learn whether the principal can commit to a long-term contract and then choose the right model.

1.5.2 Sketch of the proof

In this subsection we sketch the proof of Theorem 1, and the details will be in Appendix A. The proof uses the ergodic theory of discrete-time Markov processes. For details of the theory and relevant definitions, we refer the reader to Meyn and Tweedie (2009) (henceforth MT). Assumptions 1-5 are maintained throughout.

The probability distribution of $w_c(X, Z)$, the state in the next period, is uniquely determined by the state w in the current period for $w \in [0, \bar{w}] \setminus W_I$. To emphasize the dependence of w_c on w , we will use the notation $w_c(x, z; w)$ from now on. Define

$$P(w, W) = \text{Prob}(w_c(X, Z; w) \in W)$$

for $w \in [0, \bar{w}] \setminus W_I$ and Borel subsets $W \subset [0, \bar{w}]$. Then when $w_t \in [0, \bar{w}] \setminus W_I$, $P(w_t, \cdot)$ is the probability distribution of w_{t+1} conditional on w_t in the optimal equilibrium.

The set W_I consists of countably many connected components, or disjoint open sets. If w is in a connected component (w', w'') of W_I , then define

$$P(w, W) = \frac{w'' - w}{w'' - w'} P(w', W) + \frac{w - w'}{w'' - w'} P(w'', W)$$

for all Borel subsets W of $[0, \bar{w}]$. In other words, $P(w, \cdot)$ is defined by convex combination of P at the end points w' and w'' for $w \in (w', w'')$. We have thus defined the Markov transition probability $P(w, \cdot)$ for all $w \in [0, \bar{w}]$. The transition probability can be iterated to give $P^n(w, \cdot)$ such that $P^n(w_t, \cdot)$ is the probability distribution of w_{t+n} conditional on w_t in the optimal equilibrium. The transition probability also induces an operator T on bounded measurable functions:

$$Ta(w) = \int a(w') P(w, dw'). \quad (1.5.1)$$

In words, $Ta(w)$ is the expectation of $Ta(w_1)$ if $w_0 = w$.

Notice that when $w_t \in W_I$, $P(w_t, \cdot)$ is not necessarily the probability distribution of w_{t+1} conditional on w_t . However, we have seen that when $w_0 \notin W_I$, $w_t \notin W_I$ almost surely for all t , so if we can show that the Markov process converges to a unique invariant distribution π , we will know that the distribution of w_t converges to π for $w_0 \notin W_I$. Then a separate argument will be given for $w_0 \in W_I$.

The proof is based on the following result from the ergodic theory of Markov processes (MT, Theorem 13.0.1):

Theorem 2. If the Markov process is an aperiodic positive Harris chain, then there exists a unique invariant probability measure π such that for every initial state $w \in S$, $\sup_{A \in \mathcal{B}(S)} |P^n(w, A) - \pi(A)| \rightarrow 0$ as $n \rightarrow \infty$.

Here S is the state space of the Markov process, which is the compact set $[0, \bar{w}]$ in our model. A Markov process on a compact state space is a positive Harris chain if it is a T -chain with a reachable state w^* (MT, Theorem 18.3.2). A state w^* is reachable if for every neighborhood O of w^* if $\sum_n P^n(y, O) > 0$ for all $y \in S$. A Feller and ψ -irreducible Markov process is a T -chain if it has an open small set with positive ψ measure, or the

support of ψ has non-empty interior (MT, Propositions 6.2.5 and 6.2.8). A set C is called *small* if there exists a non-trivial measure ν on $\mathcal{B}(S)$ such that for some $m > 0$, $P^m(x, B) \geq \nu(B)$ for all $x \in C$ and $B \in \mathcal{B}(X)$.

The rest of the proof is to check that all the conditions in the theorem are satisfied by the transition probability P constructed above. The key steps are the following two lemmas:

Lemma 2. T maps continuous functions to continuous functions: for every continuous function a , $Ta(w_n) \rightarrow Ta(w)$ if $w_n \rightarrow w$.

Lemma 3. In any non-trivial optimal equilibrium, the state w_t reaches every interval $[w', w''] \subset [0, \bar{w}]$ with positive probability, provided that $V'(w') > V'(w'')$. More precisely, $\text{Prob}(w_t \in [w', w''] \text{ for some } t > 0) > 0$ if $V'(w') > V'(w'')$. In addition, the state reaches zero and \bar{w} with positive probability.

The first lemma shows that our Markov process has the Feller property: its operator T preserves continuity. This result is true by construction when w_n is a sequence in W_I . It is also easy to prove when w_n is in $[0, \bar{w}] \setminus W_I$ and the Lagrange multipliers κ_B and κ_G in Eq. (1.3.6) are not both zero at w , since in this case $w_c(x, z; w_n)$ converges to $w_c(x, z; w)$ for almost all (x, z) , which can be proved using convergence of Lagrange multipliers. The proof of the lemma when $\kappa_B = \kappa_G = 0$ is more involved. In this case, $w_c(x, z; w)$ is not characterized by the first-order condition Eq. (1.3.6). The proof uses the weak* convergence technique similar to the proof of Theorem 5 of APS, since in infinite dimensions, sequential compactness only holds for the weak* topology.

The second lemma is about the irreducibility and aperiodicity of the Markov process. It states that starting from every $w_0 > 0$, every extreme point of the frontier can be reached. When the agent is not protected by limited liability, the state zero is not absorbing, so all extreme points of the frontier can also be reached from $w_0 = 0$. When the agent is protected by limited liability, the state zero is absorbing, and the lemma asserts that from every initial state, there is positive probability that the process is absorbed by zero; in other words, the relationship is terminated. This lemma has implications about the support of the limit distribution. When the agent is not protected by limited liability, the support of π must be the set of all extreme points $[0, \bar{w}] \setminus W_I$ with a point mass at zero. When the agent is protected by limited liability, the support of π must contain

zero and assign $w = 0$ a positive mass. In fact, Theorem 1 eventually implies that the unit mass at zero is the unique invariant distribution of the Markov process.

The conditions in Theorem 2 can easily be verified using the two lemmas. It implies that the Markov process converges to the unique invariant distribution π . To prove Theorem 1, it remains to show that it is true when $w_0 \in W_I$. The idea is as follows. Since the agent's utility u is strictly concave, staying in a linear segment of V forever means that the payment b_t to the agent stays the same at all on-path histories, which implies that there is no incentive. This is impossible unless the principal and the agent repeats the static equilibrium and the agent is protected by limited liability, in which case the steady state has been reached. Once the state leaves W_I , it never comes back, and the convergence result for $w \in [0, \bar{w}] \setminus W_I$ implies that the process converges to π once it leaves W_I .

1.5.3 Long-run effects of the initial state

Although Assumptions 1-5 are not necessary for the convergence result in Theorem 1, there are some obvious examples that violate at least one of the assumptions and do not have the convergence result. In this subsection, we consider the case where the agent is protected by limited liability, but is risk-neutral at sufficiently high levels of consumption: $u(b)$ is strictly concave for $b \leq b_0$ and $u(b) = u_0 + b$ for $b \geq b_0$, where b_0 is a non-negative number and u_0 is a constant. For simplicity, we also assume that $y > c_B$, so it is efficient to exert effort in the bad environment.

The model in this subsection is a generalization of the model considered in Fong and Li (2012). They do not consider the fluctuating environment, and assume that the agent's is risk neutral at all levels of consumption (corresponding to $b_0 = 0$) and the output X_t is binary. It turns out that the model has an additional structure when X_t has continuous density, and allows us to have a slightly stronger result on the long-run dynamics.

If the payment is always above b_0 , then the model becomes a standard model of relational contract with transferable utility. In particular, one should expect the existence of stationary equilibria. In this setting, an equilibrium is called stationary if it always stays on the linear segment of the frontier V with slope -1 on the equilibrium path. Since $y > c_B$, we will focus on stationary equilibria in which the agent exerts effort in

both environments. Suppose that stationary equilibria exist and let w_m be the agent's minimum payoff in any stationary equilibrium. Using Eq. (1.3.2) to rewrite the objective function in Eq. (1.3.1), we can reduce the program Eqs. (1.3.1)(1.3.3) to the following for stationary equilibria:

$$V(w) = \max_{h \geq u(b_0), w_c \geq w_m} (1 - \delta)(y - E[c_\Theta]) + \delta V(w) - (1 - \delta)(w - u_0) \quad (1.5.2)$$

$$s.t. \quad (1 - \delta)(h - E[c_\Theta]) + \delta E[w_c(X, Z)|e = (1, 1)] = w; \quad (1.5.3)$$

$$\delta \int w_c(x, z)[g(x, 1) - g(x, 0)]f(z, \theta)dx dz \geq (1 - \delta)c_\theta. \quad (1.5.4)$$

Notice that the objective function is independent of the choice variables. Therefore, $V(w) = y - \sum_\theta \mu_\theta c_\theta - w + u_0$ as long as the constraints Eqs. (1.5.3) and (1.5.4) can be satisfied for some choice of $h \geq u(b_0)$ and $w_c \geq w_e$. In particular, if there exists a stationary equilibrium with payoff pair $(w, V(w))$ and $V(w) > 0$, then there exists stationary equilibrium that gives the agent payoff $(w + \epsilon)$ for $\epsilon \in (0, V(w)]$ since the principal can provide the additional payoff to the agent by increasing the up-front payment. Therefore, when a stationary equilibrium exists, it exists for all $w \in [w_m, \bar{w}]$, and $\bar{w} = \bar{s} \equiv y - E[c_\Theta] + u_0$, the maximum surplus that can be created in equilibrium. The following proposition characterizes the lower bound w_m :

Proposition 7. If there exists a stationary equilibrium, then V is linear on $[w_m, \bar{w}]$ where $\bar{w} = \bar{s}$, and $w_m = \omega(w_m)$, where the function ω is defined by the following program:

$$\omega(w) = \min_{w_c(x, z) \in [w, \bar{s}]} (1 - \delta)(u(b_0) - E[c_\Theta]) + \delta E[w_c(X, Z)|e = (1, 1)]; \quad (1.5.5)$$

$$s.t. \quad \delta \int w_c(x, z)[g(x, 1) - g(x, 0)]f(z, \theta)dx dz \geq (1 - \delta)c_\theta. \quad (1.5.6)$$

Moreover, $V'(w) > -1$ for $w < w_m$. Conversely, if there exists a w_m such that $\omega(w_m) = w_m$, then a stationary equilibrium exists for all $w \in [w_m, \bar{s}]$. Finally, ω always has a fixed point when \bar{s} is sufficiently big.

The $\omega(w)$ defined in Eqs. (1.5.5)-(1.5.6) is the agent's minimum expected payoff that admits an incentive compatible continuation payoff function valued on $[w, \bar{s}]$. Clearly, the up front payment at w_m is b_0 . The proposition also asserts that if ω has a fixed point, then there exists a stationary equilibrium. Therefore, the existence of a fixed point of ω is a necessary and sufficient condition for the existence of a stationary equilibrium. In addition, the slope of the equilibrium frontier is bigger than -1 for $w < w_m$, so when w_m

exists, there exists a stationary optimal equilibrium at w if and only if w is on the linear segment of V with slope -1 . In what follows, assume that a stationary equilibrium exists.

Now the Markov process no longer has a unique steady state. In fact, once w_t enters $[w_m, \bar{w}]$, it does not come out by definition of stationary equilibrium. Therefore, the state of the Markov process may end up at zero or in $[w_m, \bar{w}]$ in the long run. The following proposition says that these are the only possibilities:

Proposition 8. For any initial state w_0 , $\text{Prob}(w_t \in \{0\} \cup [w_m, \bar{w}])$ converges to one.

The proof still uses the ergodic theorem Theorem 2. As mentioned above, the Markov process does not have a unique invariant distribution π any more, but we can change the state space to make the invariant distribution unique. In particular, we create a new state space S from $[0, \bar{w}]$ by identifying 0 and all points in $[w_m, \bar{w}]$. This trick does not spoil the Markov property as once the state enters $\{0\} \cup [w_m, \bar{w}]$ which is now a single state, it stays there. The proof of the proposition verifies that all conditions in Theorem 2 hold, and non-extreme initial values can be handled as in Theorem 1. It is easy to see that the unique invariant distribution is the unit mass at $\{0\} \cup [w_m, \bar{w}]$.

The previous proposition shows that the state cannot stay in $(0, w_m)$ forever, but it does tell us whether it converges to zero or enters the regime of stationary equilibria. This issue is addressed in the following proposition:

Proposition 9. If $w_0 \in (0, w_m)$, then w_t is absorbed by zero with positive probability, and is absorbed by $[w_m, \bar{w}]$ also with positive probability.

In general, computing absorption probability of a Markov process is non-trivial. Fortunately, in order to prove this proposition, it suffices to use a martingale property of $V'(w_t)$ related to the inverse Euler equation of Rogerson (1985).

Proof. Notice that when $w_c(x, z) \in (0, \bar{w})$ for almost all (x, z) , the first-order condition Eq. (1.3.6) implies that

$$\begin{aligned} E[V'(w_c(X, Z)) - V'(w)|\mathbf{e}] &= \sum_{\theta} \mu_{\theta} \int [V'(w_c(x, z)) - V'(w)] g(x, e_{\theta}) f(z, \theta) dx dz \\ &= - \sum_{\theta} \int \kappa_{\theta} f(z, \theta) [g(x, 1) - g(x, 0)] dx dz \\ &= 0. \end{aligned}$$

This implies that $E[V'(w_{t+1})|w_t] = V'(w_t)$. However, this martingale property is spoiled by corner solutions. Suppose that the state w_t never reaches zero. Then there are only corner solutions at \bar{w} , so $V'(w_t)$ is a bounded sub-martingale. By Doob's martingale convergence theorem, it converges a.s. and in L^1 to some k_∞ . Proposition 7 implies that w_t is absorbed by $\{0\} \cup [w_m, \bar{w}]$ with probability one, and by assumption w_t never reaches zero, so w_t is eventually absorbed by $[w_m, \bar{w}]$ and thus $k_\infty = -1$ almost surely. However, $E[k_\infty] = -1 < V'(w_0)$, violating the property of a sub-martingale. Therefore, w_t is absorbed by zero with positive probability. When the state never reaches $[w_m, \bar{w}]$, $V'(w_t)$ is a bounded super-martingale, and thus converges to some k_∞ a.s. and in L^1 . The property of a super-martingale requires that $V'(w_0) \geq E[V'(w_t)] \geq E[k_\infty] = V'(0)$ for all t . However, Proposition 2 then implies that the principal never makes any payment, and thus $w_0 = 0$, a contradiction. \square

Therefore, in this example the limit distribution of w_t will depend on the initial state. In particular, the performance of an organization currently at state zero and the performance of an organization currently in the regime of stationary equilibria will not converge. Moreover, whenever the initial state is not already zero or stationary, it reaches zero and the stationary regime both with positive probability. In other words, persistent performance difference may arise from this initial state.

The martingale convergence theorem relies on the first order condition Eq. (1.3.6). Fong and Li (2012) considers a model where the output is binary and the agent is risk averse for all positive levels of consumption. When effort is binary and $b = 0$, $w_c(x; w)$ is determined by the two constraints Eq. (1.3.2) and (1.3.3) and the first-order condition Eq. (1.3.6) plays no role. (The equilibrium frontier V is not differentiable in their paper.) The above discussion shows how their result generalizes to the case when the agent is risk neutral for b above some $b_0 \geq 0$ when the output is assumed to have probability density.

In the literature of organizational economics, many models attribute persistent performance difference to persistent uncertainty about the environment and information flow about it. PPD arises because the organization may stop learning about the persistent environment before finding the efficient scheme under full information. Examples include Chassang (2010) and Halac, Kartik, and Liu (2013). These models often only consider a single organization. When there are two or more organizations, and the uncertainty

about the environment (such as the value of a new technology) is common to all organizations, each organization can learn from one another, and the cost of doing that is very different from learning by experimenting by itself. It is not clear whether the PPD created by the stoppage of learning is robust.

The PPD created in this model is different. The environment is i.i.d. over time, so there is no persistence in the environment. However, the future equilibrium path depends on the state variable w_t , which depends on history. The dynamics is driven by the principal's keeping her promise about the agent's future payoff, and the evolution of w_t in one organization does not affect the evolution of w_t in another. Therefore, if an organization is in a bad state (close to zero), it cannot improve its performance by observing the superior performance of another organization and learning from it. In industries with high information flow between participants, the evolution of w_t may be an important factor in persistent performance difference.

1.6 Extension

The ergodic theory developed in the previous section applies to more general settings than the model with binary effort and binary environment. In this section, we demonstrate that the type of convergence in Theorem 1 also holds in a setting that allows for continuous effort and any finite number of states. However, dealing with the limited liability constraint is more difficult in this more general setting, and will be left for future research.

Keep timing as the same as the baseline model. Now assume that the environment Θ_t can take a value on a general finite set, and the agent's effort e_t is a real number in $[0, 1]$. The cost of the agent's effort is $c_\theta(e)$, which is assumed to be twice continuously differentiable and convex, for every θ . Moreover, $c'_\theta(0) > 0$ for every θ . At the end of each period, a signal Z_t with domain Ω is publicly observed, where Ω is a connected open subset of a finite dimensional Euclidean space.¹¹ The distribution of Z_t only depends on e_t and Θ_t and is linear in e_t . Specifically, the probability density function of Z_t conditional on e_t and θ_t is $e_t f_1(z, \theta_t) + (1 - e_t) f_0(z, \theta_t)$, where for each θ , $f_1(\cdot, \theta)$ and $f_0(\cdot, \theta)$ are analytic functions on Ω . The assumption of linearity in e_t makes it easy to

¹¹More generally, Ω can be a finite-dimensional connected analytic manifold, such as a Riemann surface.

verify the validity of the first-order approach. The assumption that $f_1(\cdot, \theta)$ and $f_0(\cdot, \theta)$ are analytic is restrictive, and is made to facilitate the prove the bang-bang property of the optimal equilibrium.

The principal's payoff at Period t is $R(Z_t) - b_t$ for some measurable function R on Ω , and the agent's payoff at Period t is $u(b_t) - E[c_\Theta(e_t)]$. As mentioned above, the agent is not protected by limited liability, so b_t can be any real number. Notice that there is no X_t in this model: the agent's output is $R(Z_t)$. Normalize the payoff function so that $u(0) = 0$, and $E[R(Z_t)|e = 0] = 0$ in every environment. Let

$$y_\theta = E[R(Z_t)|e = 1, \theta] = \int R(z)f_1(z, \theta)dz.$$

Then the expected output in Environment θ when the agent exerts effort e_θ is $e_\theta y_\theta$.

The equilibrium frontier V and its domain $[0, \bar{w}]$ is defined as before. Consider the following program:

$$\mathcal{FV}(w) = \max_{w_c, h, e} -(1 - \delta)E[y_\Theta e_\Theta - u^{-1}(h)] + \delta E[V(w_c(Z))|e]; \quad (1.6.1)$$

$$s.t. \quad (1 - \delta)E[h - c_\Theta(e_\Theta)] + \delta E[w_c(Z)|e] = w; \quad (1.6.2)$$

$$\delta \int w_c(z)[f_1(z, \theta) - f_0(z, \theta)]dz = (1 - \delta)c'_\theta(e_\theta). \text{ for every } \theta. \quad (1.6.3)$$

This is the analogue of the program Eqs. (1.3.1)-(1.3.3). One difference is that the effort level e is now a choice variable. Since the agent's continuation payoff is linear in e and his cost of effort is convex, his IC constraint is given by the first-order condition Eq. (1.6.3). The equilibrium frontier V is the "concavification" of \mathcal{FV} .

Proposition 1 still applies. In particular, V is monotonically decreasing and continuously differentiable. Now both the objective function and the constraints are continuously differentiable in the choice variables, so there exist Lagrange multipliers κ_θ for (1.6.3) so that at optimality,

$$V'(w_c(z; w)) - V'(w) = -\frac{\sum_\theta \kappa_\theta [f_1(z, \theta) - f_0(z, \theta)]}{\sum_\theta \mu_\theta [e_\theta(w)f_1(z, \theta) + (1 - e_\theta(w))f_0(z, \theta)]}. \quad (1.6.4)$$

Notice that the Lagrange multiplier for Eq. (1.6.2) is $-V'(w)$ by Proposition 1. Since the objective function in Eq. (1.6.1) is not concave in the choice variables, the first order

condition may not be sufficient for optimality. Fortunately, the sufficiency is never used in the proof of Theorem 1.

It is still assumed that $\bar{w} > 0$. As before, let $W_I = \{w \in (0, \bar{w}) : V \text{ is linear on a neighborhood of } w\}$, and assume that Assumptions 3 and 4 hold. The monotone likelihood ratio assumption and Assumption 2 are replaced by the assumption that $f_1(\cdot, \theta)$ and $f_0(\cdot, \theta)$ are analytic. Finally, u is still assumed to be strictly concave, and $f_1(\cdot, \theta)/f_0(\cdot, \theta)$ is bounded away from zero and infinity on Ω for all θ .

Proposition 10. Under the above assumptions, $\text{Prob}(w_{t+1} \in W_I | w_t \notin W_I) = 0$. Moreover, there exists a unique probability distribution π on $[0, \bar{w}]$ such that in every optimal equilibrium, $\sup_{A \in \mathcal{B}([0, \bar{w}])} |\text{Prob}(w_t \in A) - \pi(A)| \rightarrow 0$ as $t \rightarrow \infty$. Moreover, the support of π is $[0, \bar{w}] \setminus W_I$, $\pi(\{0\}) > 0$, and $\pi(\{\bar{w}\}) > 0$.

This result is the combination of Proposition 5 and Theorem 1 in this new setting. In fact, the proof of this result simply checks that the proof of the two old results applies in this setting with small modifications. For example, the proof of Proposition 5 relies on the strict monotonicity of the likelihood ratio functions, and the proof of the bang-bang property in this setting relies on fact that zeros of an analytic function are isolated, which is closer to the proof of the bang-bang property in APS.

Therefore, we have seen that the mathematical tool of Markov processes applies to a more general setting than the binary-effort-binary-environment baseline model, and the optimal equilibria of this more general model still have the ergodic property of Theorem 1. However, this generalization is not without price. First, dealing with limited liability is more difficult because the differentiability of V cannot be established in the straight forward way of Proposition 1. Secondly, the assumption that the probability density functions are analytic is very restrictive, and rules out many potentially interesting density functions. Thirdly, results in Section 4 that relate the characterization of the equilibrium frontier to empirical evidence are lost.

1.7 Conclusion

In a principal-agent relationship, optimal dynamic response to output and environment gives rise to rich dynamics when the agent is risk-averse and/or protected by limited

liability. The optimal dynamic response generates empirical implications that are different from the conventional wisdom and sometimes fits data better. In the long-run, characteristics of equilibrium may or may not converge to a fixed distribution, which opens way towards persistent performance difference. Interesting questions can also be asked about finite-time evolution of the agent's effort level, but will be left for future research.

1.8 Appendices

1.8.1 Appendix A: proofs

This appendix collects proofs of results that are not included in the main text.

Proof of Proposition 2. We first prove the third assertion. Suppose in equilibrium the agent exerts effort in environment θ . Consider the following program:

$$\begin{aligned} \min_{w_c} \quad & \int w_c(x, z)g(x, 0)f(z, \theta)dx dz \\ \text{s.t.} \quad & \int w_c(x, z)[g(x, 1) - g(x, 0)]f(z, \theta)dx dz \geq \frac{1 - \delta}{\delta}c_\theta. \end{aligned}$$

Under Assumption 1, the strong duality holds for this linear program, so the optimal solution minimizes the Lagrangian $\int w_c(x, z)[g(x, 0) - \lambda(g(x, 1) - g(x, 0))]f(z, \theta)dx dz$ for some $\lambda \geq 0$. Clearly, the solution involves setting $w_c(x, z) = \bar{w}$ for x above some threshold and $w_c(x, z) = 0$ for x below that threshold. The value of the program is some positive number \underline{w}_θ . Now consider the program Eqs. (1.3.1)-(1.3.3). Since the agent can always choose not to exert effort in either environment and $h \geq 0$, his payoff is at least

$$\delta\mu_\theta \int w_c(x, z)g(x, 0)f(z, \theta)dx dz.$$

Since w_c satisfies Eq. (1.3.3) with $e_\theta = 1$, the above expression is at least $\delta\mu_\theta\underline{w}_\theta$. Therefore, when $w < \delta\mu_\theta\underline{w}_\theta$, it is impossible to induce effort in environment θ . Consequently, $\mathcal{F}_e V(w) = -\infty$ for $e \neq (0, 0)$ and $w < w_s \equiv \min\{\delta\mu_B w_B, \delta\mu_G w_G\}$. Therefore, $V(0) = 0$, and for $w < w_s$, the optimal equilibrium either involves the two players' playing the static equilibrium in the first period, or involves a public randomization in the first period. In the latter case, V is locally linear on w . In the former case, it is optimal to set $w_c(x, z)$ independent of (x, z) due to the concavity of V , so $w_c(x, z) = \delta^{-1}w$ and $V(w) = \delta V(\delta^{-1}w)$, which means that V is linear on $[0, \delta^{-1}w]$. We conclude that V is linear on $[0, w_s]$.

Now we turn to the second assertion of the proposition. In order to prove the strong duality, we need to show that there exists some (h, w_c) that satisfies Eq. (1.3.2) and Eq. (1.3.3) with strict inequality. It is easy to see that such an (h, w_c) does not exist if and

only if

$$w = \min_{w_c} -(1 - \delta) \sum_{\theta} \mu_{\theta} c_{\theta} e_{\theta} + \delta \int w_c(x, z) g(x, e_{\theta}) f(z, \theta) dx dz \quad (1.8.1)$$

$$s.t. \quad (2e_{\theta} - 1)\delta \int w_c(x, z) [g(x, 1) - g(x, 0)] f(z, \theta) dx dz \geq (1 - \delta) c_{\theta}. \quad (1.8.2)$$

(If w is bigger than the value of the above program, choose w_c as in the program, modify it a bit so that the IC constraints are satisfied with strict inequality¹² and then choose a positive h .) For greater w , the strong duality holds, which leads to the following first-order conditions:

$$\frac{1}{u'(u^{-1}(h))} = \gamma; \quad (1.8.3)$$

$$V'(w_c(x, z)) = -\gamma - \frac{\kappa_G f(z, G) + \kappa_B f(z, B)}{\mu_G g(x, e_G) f(z, G) + \mu_B g(x, e_B) f(z, B)} [g(x, 1) - g(x, 0)]. \quad (1.8.4)$$

Here γ is the Lagrange multiplier for Eq. (1.3.2), and κ_{θ} is the Lagrange multiplier for Eq. (1.3.3). When V is not differentiable, the second condition means that the right hand side is between the left and right derivatives of V at $w_c(x, z)$.

The solution to the program Eqs. (1.8.1)-(1.8.2) involves setting each $w_c(x, z)$ to be either 0 or \bar{w} . Either way $V(w_c(x, z)) = 0$. This implies that $\mathbf{e} \neq (1, 0)$, since in that case the IC constraint for $\theta = B$ must be binding, and asking the agent to exert effort in the bad environment does not change either player's continuation payoff but does increase the principal's payoff in the current period. Moreover, the value of the program is zero when $e_B = e_G = 0$. We have seen that Eq. (1.3.6) holds when $w = 0$. It remains to consider the possibility that $\mathbf{e} = (1, 1)$ or $(0, 1)$. In the latter case, it is easy to see that the IC constraint for $\theta = B$ in the program Eqs. (1.8.1)-(1.8.2) is not binding.

Due to Assumption 1, with positive probability $w_c(X, Z) = 0$ and $g(X, 1) > g(X, 0)$. Now when the agent's expected payoff increases from w to $w + \epsilon$, the principal can raise $w_c(x, z)$ with $g(x, 1) > g(x, 0)$ to accommodate this increase without spoiling the agent's IC constraints. For sufficiently small ϵ , the increase in $w_c(x, z)$ can be bounded by w_s . This results an increase in the principal's payoff by $V'(0)\epsilon$. However, this choice of the agent's continuation payoff is not optimal, as under Assumption 2, the support of $V'(w_c(x, z))$ should be a connected by Eq. (1.8.4). Therefore, $V(w + \epsilon) > V(w) + V'(0)\epsilon$,

¹² This is possible under Assumption 1.

contradicting the concavity of V . To sum up, $V_e(w)$ cannot be optimal when w solves the program Eqs. (1.8.1)-(1.8.2). Consequently, the first-order conditions Eqs. (1.8.3)-(1.8.4) always hold for some γ , κ_B , and κ_G .

Now we prove the differentiability of V . Suppose that in the program Eqs. (1.3.1)-(1.3.3) at $w \in (0, \bar{w})$, the optimal choice of h is strictly positive. Then the proof of Proposition 1 applies and shows that V is differentiable at w and $V'(w) = -1/u'(u^{-1}(h(w)))$. When the optimal h is zero, Eq. (1.3.5) still holds for positive ϵ , and taking the right derivative of both sides at $\epsilon = 0$ yields that $V'(w+) \geq -1/u'(0)$. Therefore, V is differentiable whenever $V'(w+) \geq -1/u'(0)$. Let $w_d = \inf\{w : V' \text{ is differentiable on } (w, \bar{w})\}$.

Suppose that $w_d > 0$. If $V'(w_d+) < -1/u'(0)$, then V is differentiable at w_d and the right derivative of V is continuous at w_d . This implies that the right derivative of V is smaller than $-1/u'(0)$ in a neighborhood of w_d , and thus V is differentiable on this neighborhood, a contradiction. Therefore, $V'(w_d+) \geq -1/u'(0)$, and thus $h(w) = 0$ for all $w \leq w_d$. We show that V is differentiable at all $w \in (\delta w_d, w_d]$ to derive a contradiction.

Consider a $w \in (\delta w_d, w_d]$. If a public randomization is strictly optimal at w , V is locally linear at w and thus differentiable. It remains to consider the case where $V(w) = \mathcal{F}_e V(w)$ for some e . By Eq. (1.3.2), $w_c(X, Z) > w_d + \frac{1}{2}\xi$ with positive probability, where $\xi = \delta^{-1}w - w_d$. By Eq. (1.8.4) and the continuity of the likelihood ratios, the set of (x, z) such that $w_c(x, z) > w_d + \frac{1}{2}\xi$ is open. Therefore, there exists $A = [x_1, x_2] \times [z_1, z_2]$ such that a) $w_c(x, z) > w_d + \frac{1}{2}\xi$ for $(x, z) \in A$; b) $g(x, e)f(z, \theta) > 0$ on A for all e and θ ; and c) $g(x_1, 1) - g(x_1, 0)$ has the same sign as $g(x_2, 1) - g(x_2, 0)$. Choose $x_m \in (x_1, x_2)$ and $z_m \in (z_1, z_2)$, and let

$$D_{e\theta, i} \equiv \int_{B_i} g(x, e)f(z, \theta) dx dz,$$

for $e \in \{0, 1\}$, $\theta \in \{B, G\}$, and $i \in \{1, 2, 3, 4\}$, where $B_1 = [x_1, x_m] \times [z_1, z_m]$, $B_2 = [x_1, x_m] \times [z_m, z_2]$, $B_3 = (x_m, x_2] \times [z_1, z_m]$, and $B_4 = (x_m, x_2] \times [z_m, z_2]$.

For $\epsilon = (\epsilon_1, \epsilon_2, \epsilon_3)$, define $\tilde{w}_c(x, z; \epsilon)$ by

$$\tilde{w}_c(x, z; \epsilon) = \begin{cases} w_c(x, z) + \epsilon_i & , \text{ if } (x, z) \in B_i, i = 1, 2, 3; \\ w_c(x, z) & , \text{ otherwise.} \end{cases}$$

Replacing w_c with \bar{w}_c in the program Eqs. (1.3.1)-(1.3.3), we see that the left hand sides of Eq. (1.3.2), (1.3.3) with $\theta = B$, and (1.3.3) with $\theta = G$ are linear functions of ϵ , with Jacobian matrix

$$J = \delta \begin{pmatrix} \sum_{\theta} \mu_{\theta} D_{e_{\theta},1} & \sum_{\theta} \mu_{\theta} D_{e_{\theta},2} & \sum_{\theta} \mu_{\theta} D_{e_{\theta},3} \\ (2e_B - 1)(D_{1B,1} - D_{0B,1}) & (2e_B - 1)(D_{1B,2} - D_{0B,2}) & (2e_B - 1)(D_{1B,3} - D_{0B,3}) \\ (2e_G - 1)(D_{1G,1} - D_{0G,1}) & (2e_G - 1)(D_{1G,2} - D_{0G,2}) & (2e_G - 1)(D_{1G,3} - D_{0G,3}) \end{pmatrix}.$$

Notice that

$$\frac{D_{1B,1} - D_{0B,1}}{D_{1G,1} - D_{0G,1}} = \frac{D_{1B,3} - D_{0B,3}}{D_{1G,3} - D_{0G,3}} = \frac{\int_{z_1}^{z_m} f(z, B) dz}{\int_{z_1}^{z_m} f(z, G) dz} > \frac{\int_{z_m}^{z_2} f(z, B) dz}{\int_{z_m}^{z_2} f(z, G) dz} = \frac{D_{1B,2} - D_{0B,2}}{D_{1G,2} - D_{0G,2}},$$

and that for $e \in \{0, 1\}$,

$$\begin{aligned} & (D_{1G,1} - D_{0G,1}) \int_{x_m}^{x_2} g(x, i) dx - (D_{1G,3} - D_{0G,3}) \int_{x_1}^{x_m} g(x, i) dx \\ &= \left[\int_{x_1}^{x_m} g(x, 1) dx \int_{x_m}^{x_2} g(x, 0) dx - \int_{x_m}^{x_2} g(x, 1) dx \int_{x_1}^{x_m} g(x, 0) dx \right] \int_{z_1}^{z_m} f(z, G) dz < 0. \end{aligned}$$

Compute the determinant of J by first subtracting $(2e_B - 1)(D_{1B,1} - D_{0B,1}) / [(2e_G - 1)(D_{1G,1} - D_{0G,1})]$ times the third row from the second row and then expanding by the second row. We see that the sign of $\det J$ is the same as the sign of $-(2e_B - 1)(2e_G - 1)(D_{1G,2} - D_{0G,2})$. In particular, $\det J \neq 0$. Therefore, for sufficiently small η , a change of the agent's expected payoff by η can be accommodated by changes in $(\epsilon_1, \epsilon_2, \epsilon_3)$ without changing the left hand sides of the IC constraints, and furthermore, $(\epsilon_1, \epsilon_2, \epsilon_3)$ are linear functions of η . When η is sufficiently small, $\max_i |\epsilon_i(\eta)| < \xi/2$, so the principal's payoff is differentiable in the ϵ_i 's and thus in η . It is also easy to see that the derivative at $\eta = 0$ is the $-\gamma$ in Eq. (1.8.4). Therefore, the concave function V is bounded from below by a differentiable function in a neighborhood of w and the bound is tight at w , so V is differentiable at w and the derivative is $-\gamma$. This holds for all $w \in (\delta w_d, w_d]$, contradicting the definition of w_d . Therefore, V is differentiable on $(0, \bar{w})$, and the derivative is $-\gamma$ in Eq. (1.8.4) when $V(w) = V_e(w)$. Therefore, Eq. (1.3.6) holds.

Proof of Proposition 3. First consider the case where $e_B = e_G = 0$. Ignore the agent's IC constraints in both environments for now. Then the principal maximizes $E[V(w_c(X, Z))]$ with the only constraint being on $E[w_c(X, Z)]$. Since V is concave, it

is always optimal to make $w_c(x, z)$ independent of both x and z . Clearly, this choice satisfies the agent's IC constraints.

Next consider the case $(e_B, e_G) = (0, 1)$. Ignore the incentive-compatibility constraint for $\theta = B$ for now. Then $\kappa_B = 0$ in Eq. (1.3.6). The monotone likelihood ratio property of f implies that the right hand side of Eq. (1.3.6) is weakly increasing in z for low output, and weakly decreasing in z for high output, and the sign of the right hand side of Eq. (1.3.6) is the opposite to the sign of $g(x, 1) - g(x, 0)$. Therefore, $V'(w_c(x, z)) \geq V'(w)$ and $V'(w_c(x, z))$ is weakly increasing in z for low output x ; $V'(w_c(x, z)) \leq V'(w)$ and $V'(w_c(x, z))$ is weakly decreasing in z for high output x . Since V' is weakly decreasing, this implies the desired result.

For each high output x , since $F(\cdot, G)$ first-order-stochastically-dominates $F(\cdot, B)$,

$$\int w_c(x, z)f(z, B)dz \leq \int w_c(x, z)f(z, G)dz,$$

and for each low output x , the reverse inequality holds. This implies that

$$\int w_c(x, z)(g(x, 1) - g(x, 0))f(z, B)dx dz \leq \int w_c(x, z)(g(x, 1) - g(x, 0))f(z, G)dx dz.$$

Hence, the agent's incentive is weaker in the bad environment, and his incentive-compatibility constraint for $\theta = B$ is automatically satisfied.

The monotonicity of w_c in z is trivial only if $\kappa_G = 0$. In this case, $V'(w_c(x, z)) = V'(w)$ for all x and z , so V is linear on the convex hull of the union of support of w_c and $\{w\}$.

Finally, consider the case where $e_B = e_G = 1$. Then clearly $\kappa_B \geq 0$ and $\kappa_G \geq 0$, so the first assertion of the proposition follows. When $\kappa_G = 0$, the argument used in the previous case implies the desired monotonicity of $w_c(x, z)$ in z , and the monotonicity is trivial if and only if $\kappa_B = \kappa_G = 0$.

When $\kappa_G > 0$ the monotonicity of the right hand side of Eq. (1.3.6) on z depends on the sign of $\mu_B \kappa_G - \mu_G \kappa_B$. When $\mu_B \kappa_G > \mu_G \kappa_B$, the right hand side of Eq. (1.3.6) is strictly increasing in z for low output x and strictly decreasing in z for high output x . Therefore, the optimal $w_c(x, z)$ is strictly decreasing in z for low output x and strictly increasing in z for high output x . The argument in the previous case then implies that the agent's incentive is strictly weaker in the bad environment, which cannot be true when

the incentive-compatibility constraint is binding for $\theta = G$. Therefore, $\mu_B \kappa_G \leq \mu_G \kappa_B$. This gives the desired monotonicity result. Clearly, the monotonicity result cannot be trivial in this case, as that would imply that the agent's incentive is independent of the environment, contradicting to the assumption that his IC constraint is binding for $\theta = G$.

Proof of Corollary 1. Propositions 1 and 2 allows us to replace $V'(w_c(x, z))$ in Eq. (1.3.6) by $-1/u'(b_{t+1})$, which is a decreasing and linear function of b_{t+1} when c_2 is quadratic. Since $\mu_B \kappa_G \leq \mu_G \kappa_B$, the third term on the left hand side of Eq. (1.3.6) has decreasing differences in x and z . Hence, b_{t+1} has decreasing differences in X_t and Z_t .

Proof of Proposition 4. First consider the case where $V'(w) < 0$. Let \tilde{V} be the concavification of $\max_{e \neq (0,0)} \mathcal{F}_e V$, the principal's maximum payoff when the agent is not allowed to shirk in the first period. Since $(0, 0)$ is optimal at w , the agent's optimal continuation payoff can be chosen as independent of (x, z) and will be denoted by w' . Then

$$V(w) = -(1 - \delta)u^{-1}(h) + \delta V(w') \leq \delta V(w').$$

Since $V'(\tilde{w}) \leq V'(w) < 0$ for all $\tilde{w} \geq w$, the above inequality implies that $w' < w$. On the other hand, Eq. (1.3.6) implies that $V'(w) = V'(w')$. Therefore, V is linear on $[w', w]$. There exists $\epsilon > 0$ so that for $\tilde{w} \in (w, w + \epsilon)$, the principal receives payoff at least $V(w) + V'(w)(\tilde{w} - w)$ by inducing zero effort in the first period and giving the agent the same first-period payment and a continuation payoff of $w' + \delta^{-1}(\tilde{w} - w)$. (In fact, ϵ can be chosen as $\delta(w - w')$.) Since \tilde{V} is concave and weakly below the line through $(w', V(w'))$ and $(w, V(w))$ at w' and at w , it is weakly below the line on $[0, \tilde{w}]$. Therefore, zero effort continues to be optimal on $[w, w + \epsilon)$. This argument can be used for all $w < \tilde{w}$, so

$$\sup\{\tilde{w} \in [w, \bar{w}] : (0, 0) \text{ is optimal on } [w, \tilde{w}]\} = \bar{w}.$$

This proves the first assertion of the proposition.

Next consider the case where the agent is protected by limited liability and $V'(w) > 0$. Then the payment in the first period is zero, the agent's continuation value can be chosen as a constant, namely $\delta^{-1}w$, and $V(w) = \delta V(\delta^{-1}w)$. This means that $(0, 0)$, $(w, V(w))$

and $(\delta^{-1}w, V(\delta^{-1}(w)))$ are on the same line. It is then easy to see that $(0, 0)$ is also optimal for all $\bar{w} \in [0, w]$.

Proof of Proposition 5. First consider the case where $\max\{\kappa_B, \kappa_G\} > 0$. If $\mu_B \kappa_G = \mu_G \kappa_B$, then $e_B = e_G = 1$ and by Eq. (1.3.6), $V'(w_c(x, z))$ is independent of z and strictly decreasing in x . Since $V'(W_I)$ is a countable set, $\{(x, z) : w_c(x, z) \in W_I\}$ has measure zero. Next consider the case where $\kappa_B = \kappa_G = 0$. We prove the following result:

Lemma 4. If $\kappa_B = \kappa_G = 0$ for some $w \notin W_I$, then $(e_B, e_G) \neq (1, 0)$, w is at the boundary of a connected component (w', w'') of W_I , $w'' - w' \geq (1 - \delta)\delta^{-1}c_G$, $w_c(x, z) \in \{w', w''\}$ for almost all (x, z) . Finally, if both w_c and \tilde{w}_c are optimal, then $w_c(x, z) = \tilde{w}_c(x, z)$ almost everywhere.

Proof. We only need to consider the non-trivial case where $(e_B, e_G) \neq (0, 0)$. Since $\kappa_B = \kappa_G = 0$, $V'(w_c(x, z)) = V'(w)$ for all x and z . However, w_c cannot be identically equal to w due to the incentive-compatibility constraint, so w is at the boundary of a connected component (w', w'') of W_I . Notice that V is linear on $[w', w'']$, so the analysis of the risk-neutral benchmark applies here and implies that the principal's payoff does not depend on the choice w_c as long as the IC constraints can be satisfied. If $(e_B, e_G) = (1, 0)$, then $\delta \int (w'' - w')[g(x, 1) - g(x, 0)]^+ dx \geq (1 - \delta)c_B$, and thus both $(e_B, e_G) = (1, 1)$ and $(e_B, e_G) = (0, 0)$ can be implemented. When $(1 - \delta)y > -\delta V'(w)c_G$, $(1, 0)$ is dominated by $(1, 1)$, and when $(1 - \delta)y < -\delta V'(w)c_B$, $(1, 0)$ is dominated by $(0, 0)$.

Now assume that $w = w'$. Consider the following program:

$$w_l(w) = \min_{\tilde{w}_c} (1 - \delta)(h(w) - E[c_{\Theta} e_{\Theta}]) + \delta \sum_{\theta} \mu_{\theta} \int \tilde{w}_c(x, z) g(x, e_{\theta}) f(z, \theta) dz \quad (1.8.5)$$

$$s.t. \quad \tilde{w}_c(x, z) \in [w', w'']; \quad (1.8.6)$$

$$(2e_{\theta} - 1)\delta \int \tilde{w}_c(x, z)[g(x, 1) - g(x, 0)]f(z, \theta) dz \geq (1 - \delta)c_{\theta}. \quad (1.8.7)$$

Here $h(w)$ is the optimal choice of h in the program Eqs. (1.3.1)-(1.3.3) at w . Since w_c is feasible in the above program, $w_l(w) \leq w$. Suppose that $w_l(w) < w$. Then $(h(w), \tilde{w}_c)$ is feasible in the program Eqs. (1.3.1)-(1.3.3) for $\mathcal{F}_e V(w_l(w))$. Since both w_c and \tilde{w}_c only take values in $[w', w'']$, \tilde{w}_c gives the principal payoff $V(w) - (w - w_l(w))V'(w)$. Therefore, $V(w_l(w)) \geq \mathcal{F}_e V(w_l(w)) \geq V(w) - (w - w_l(w))V'(w)$, which implies that V is linear on $[w_l(w), w'']$, a contradiction. Therefore, $w_l(w) = w$. If Eq. (1.8.7) cannot be satisfied

with strict inequality, then clearly $w_c(x, z) = w''$ for all high output x and $w_c(x, z) = w'$ for all low output x . Suppose that Eq. (1.8.7) can hold with strict inequality. Then the linear program can be solved using the dual method, and the solution calls for $\tilde{w}_c(x, z) \in \{w', w''\}$ for almost all (x, z) . If there are two optimal choices of \tilde{w}_c that differ on a set with positive measure, their convex combination is also optimal, but will take values on (w', w'') for (x, z) with positive measure, a contradiction. Therefore, the solution is unique up to (x, z) with zero measure, and coincides with $w_c(x, z)$.

If $w = w''$, then consider the program Eqs. (1.8.5)-(1.8.7) with the min replaced by max. The same extension argument shows that the value of the new program is w , and the uniqueness and the bang-bang property of the solution follow. \square

It remains to prove the uniqueness of the optimal w_c . Assume that \tilde{w}_c is also optimal and $w_c \neq \tilde{w}_c$ on a set with positive measure. Then $(w_c + \tilde{w}_c)/2$ is also optimal and strictly dominates w_c and \tilde{w}_c unless V is linear on the convex hull of $\{w_c(x, z), \tilde{w}_c(x, z)\}$ for almost all (x, z) . Therefore, $w_c(x, z)\tilde{w}_c(x, z) \in W_I$ whenever $w_c(x, z) \neq \tilde{w}_c(x, z)$ (except perhaps for some (x, z) with zero measure). The discussion of the case $\max\{\kappa_B, \kappa_G\} > 0$ implies that this is possible only when $\kappa_B = \kappa_G = 0$ for both w_c and \tilde{w}_c . However, this possibility is ruled out by the lemma.

Proof of Lemma 2. When w is in the closure of W_I and w_n is a sequence in W_I , the construction of transition probability implies that $\sup_{A \in \mathcal{B}([0, \bar{w}])} |P(w_n, A) - P(w, A)| \rightarrow 0$ as $n \rightarrow \infty$, so $Ta(w_n) \rightarrow Ta(w)$. In what follows, assume that w_n is a sequence in $[0, \bar{w}] \setminus W_I$. Then $w \in [0, \bar{w}] \setminus W_I$ too. Suppose that $Ta(w_n)$ does not converge to $Ta(w)$. Then there is a subsequence of w_n whose image under Ta is bounded away from $Ta(w)$. Therefore, it suffices to show that every sequence in $[0, \bar{w}] \setminus W_I$ that converges to w has a subsequence whose image under Ta converges to $Ta(w)$.

Since $\mathbf{e}(w)$ is unique, $\mathbf{e}(w_n) \rightarrow \mathbf{e}(w)$. WLOG, assume that $\mathbf{e}(w_n) = \mathbf{e}(w)$ for all n . Notice that this implies that the probability distribution of (X, Z) is independent of n since its probability density is $\sum_{\theta} \mu_{\theta} g(x, e_{\theta}(w_n)) f(z, \theta)$. If $\mathbf{e}(w) = (0, 0)$, then $w_c(x, z; w)$ is independent of x and z , $V'(w_c(x, z; w)) = V'(w)$, and $V(w) = -(1 - \delta)u^{-1}(h(w)) + \delta V(w_c(x, z; w))$. Therefore, $w_c(x, z; w) \neq w$ unless $w = V(w) = 0$. This implies that w is in the closure of W_I . The same is true for all w_n . Since $w_n \notin W_I$ for any n , $w_n = w$ for all n . There is nothing to prove. In what follows, assume that $\mathbf{e}(w) \neq (0, 0)$.

Let κ_{Gn} , and κ_{Bn} be the Lagrange multipliers in Eq. (1.3.6) for w_n , and let $\kappa_{G,\infty}, \kappa_{B,\infty}$ be the Lagrange multipliers in Eq. (1.3.6) for w . Suppose that $(\kappa_{Gn}, \kappa_{Bn})$ does not converge to $(\kappa_{G,\infty}, \kappa_{B,\infty})$. Then there exists a subsequence of $\{(\kappa_{Gn}, \kappa_{Bn})\}$ that converges to some $(\tilde{\kappa}_G, \tilde{\kappa}_B) \neq (\kappa_{G,\infty}, \kappa_{B,\infty})$. WLOG assume that the subsequence is $\{(\kappa_{Gn}, \kappa_{Bn})\}$ itself.

First assume that $(\tilde{\kappa}_G, \tilde{\kappa}_B) \neq (0, 0)$. Then replacing κ_G and κ_B in Eq. (1.3.6) at w by $\tilde{\kappa}_G$ and $\tilde{\kappa}_B$ yields some $\tilde{w}_c(x, z)$. In case for some (x, z) there are many $w_c(x, z)$ that satisfy Eq. (1.3.6), pick one arbitrarily. Then unless V' is not strictly decreasing at $\tilde{w}_c(x, z)$, $w_c(x, z, w_n) \rightarrow \tilde{w}_c(x, z)$. By Proposition 5, the (x, z) such that V' is not strictly decreasing at $\tilde{w}_c(x, z)$ has Lebesgue measure zero. Therefore, $\tilde{w}_c(x, z)$ is feasible for the program Eqs. (1.3.1)-(1.3.3) at w , and $E[V(w_c(X, Z; w_n))|e(u_n)] \rightarrow E[V(\tilde{w}_c(X, Z))|e(w)]$. By the continuity of V , $\tilde{w}_c(x, z)$ is optimal. By Proposition 5, $\tilde{w}_c = w_c(\cdot, \cdot; w)$ almost everywhere, implying that $(\tilde{\kappa}_G, \tilde{\kappa}_B) = (\kappa_{G,\infty}, \kappa_{B,\infty})$, a contradiction. The above argument also shows that $w_c(x, z; w_n)$ converges to $w_c(x, z; w)$ for almost all (x, z) when $(\tilde{\kappa}_G, \tilde{\kappa}_B) \neq (0, 0)$, which implies that $Ta(w_n) \rightarrow Ta(w)$ by the dominated convergence theorem.

Next assume that $(\tilde{\kappa}_G, \tilde{\kappa}_B) = (0, 0)$. Let $[w', w''] = (V')^{-1}(V'(w'))$. If $w' = w''$, then the argument used in the previous paragraph shows that $w_c(x, z; w_n)$ converges to w' for almost all (x, z) , which violates the IC constraint. Therefore, $w'' > w'$. Since g_H/g_L is bounded, the support of $V'(w_c(X, Z; w_n))$ converges to $\{V'(w')\}$ as $n \rightarrow \infty$. Therefore, for all (x, z) such that Eq. (1.3.6) holds for all n (which includes almost all (x, z)), every convergent subsequence of $\{w_c(x, z; w_n)\}$ converges to a number in $[w', w'']$.

Denote the domain of (x, z) by Ω . Now view $L^\infty(\Omega)$ as the dual of $L^1(\Omega)$ (under the Lebesgue measure). Since $L^1(\Omega)$ is separable, the weak*-topology on $\{v \in L^\infty(\Omega) : \|v\|_{L^\infty} \leq N\}$ is metrizable and sequentially compact by the sequential version of the Banach-Alaoglu theorem. Therefore, there exists a subsequence of $(w_c(\cdot, \cdot; w_n), V(w_c(\cdot, \cdot; w_n)))$ that converges to some (\tilde{w}_c, \tilde{v}) in weak-* topology on $L^\infty(\Omega)$. WLOG assume that the subsequence is $(w_c(\cdot, \cdot; w_n), V(w_c(\cdot, \cdot; w_n)))$ itself. Let $L = \{(\tilde{w}, \tilde{v}) : \tilde{v} = V'(w')\tilde{w} + \beta\}$ be the line that contains $(w', V(w'))$ and $(w'', V(w''))$. We have seen that for almost all (x, z) , every converging subsequence of $\{w_c(x, z; w_n)\}$ converges to a number in $[w', w'']$, so for almost all (x, z) , $V(w_c(x, z; w_n)) - V'(w')w_c(x, z; w_n) - \beta \rightarrow 0$. By the dominated

convergence theorem,

$$\int [V(w_c(x, z; w_n)) - V'(w')w_c(x, z; w_n) - \beta]g(x, e)f(z, \theta)dx dz \rightarrow 0,$$

as $n \rightarrow \infty$ for all $(e, \theta) \in \{0, 1\} \times \{B, G\}$. Therefore,

$$\int [\tilde{v}(x, z) - V'(w')\tilde{w}_c(x, z) - \beta]g(x, a)f(z, \theta)dx dz = 0.$$

On the other hand, using the argument in the proof of APS's Theorem 5, we can show that $\tilde{v} \leq V \circ \tilde{w}_c$ almost everywhere. Therefore,

$$\int [V(\tilde{w}_c(x, z)) - V'(w')\tilde{w}_c(x, z) - \beta]g(x, a)f(z, \theta)dx dz \geq 0.$$

However, V is concave, so the whole equilibrium frontier is below the line L . Therefore, $\tilde{v}(x, z) = V(\tilde{w}_c(x, z))$ and $w_c(x, z) \in [w', w'']$ for almost all (x, z) . Weak*-convergence implies that $\tilde{w}_c(x, z)$ satisfies the constraints Eqs. (1.3.2) and (1.3.3), and the continuity of V further implies that \tilde{w}_c is optimal in the program of Eqs. (1.3.1)-(1.3.3) at w . By Proposition 5, $\tilde{w}_c = w_c(\cdot, \cdot; w)$ almost everywhere. However, this means that $(\kappa_{B, \infty}, \kappa_{G, \infty}) = (0, 0) = (\tilde{\kappa}_B, \tilde{\kappa}_G)$, a contradiction. Therefore, $(\kappa_{Bn}, \kappa_{Gn}) \rightarrow (0, 0) = (\kappa_{B, \infty}, \kappa_{G, \infty})$. This implies that $w = w'$ or $w = w''$. The above argument also shows that $w_c(x, z; w_n)$ converges to $w_c(x, z; w)$ in the weak*-topology.

Fix an $\epsilon > 0$. Since a is continuous on the compact space $[0, \bar{w}]$, there exists an $\eta > 0$ such that $|a(w^{(1)}) - a(w^{(2)})| < \epsilon/3$ whenever $|w^{(1)} - w^{(2)}| \leq \eta$. Also, choose $M \geq \max_{\bar{w}} a(\bar{w}) - \min_{\bar{w}} a(\bar{w})$. In fact, we can choose η and M so that

$$\left[\frac{\bar{w}\epsilon}{3M} + \eta \right] |a(w'') - a(w')| < \frac{\epsilon}{6(w'' - w')}. \quad (1.8.8)$$

We have seen that $w' \leq \liminf_{n \rightarrow \infty} w_c(x, z; w_n) \leq \limsup_{n \rightarrow \infty} w_c(x, z; w_n) \leq w''$ for almost all (x, z) . Therefore, $\text{Prob}(\sup_{n \geq m} w_c(X, Z; w_n) > w'' + \eta) \rightarrow 0$ and $\text{Prob}(\inf_{n \geq m} w_c(X, Z; w_n) < w' - \eta) \rightarrow 0$ as $m \rightarrow \infty$. Find an N_1 such that $\text{Prob}(\sup_{n \geq N_1} w_c(X, Z; w_n) > w'' + \eta) < \epsilon/(6M)$ and $\text{Prob}(\inf_{n \geq N_1} w_c(X, Z; w_n) < w' - \eta) < \epsilon/(6M)$. Let $b_n = \text{Prob}(w_c(X, Z; w_n) \geq w'' | \mathbf{e}(w_n))$, and $b = \text{Prob}(w_c(X, Z; w) \geq w'' | \mathbf{e}(w))$. Then $\text{Prob}(w_c(X, Z; w_n) \leq w' | \mathbf{e}(w_n)) = 1 - b_n$ by the bang-bang property established in Proposition 5. Then

$$Ta(w) = ba(w'') + (1 - b)a(w'), \quad (1.8.9)$$

and for $n \geq N_1$,

$$\begin{aligned} & |Ta(w_n) - b_n a(w'') - (1 - b_n)a(w')| \\ & \leq [\text{Prob}(w_c(X, Z; w_n) > w'' + \eta) + \text{Prob}(w_c(X, Z; w_n) < w' - \eta)] M + \frac{\epsilon}{3} \\ & < \frac{2\epsilon}{3}. \end{aligned}$$

The above estimate is obtained by replacing $a(w_c(x, z; w_n))$ by $a(w'')$ when $w_c(x, z; w_n) \geq w''$ and replacing $a(w_c(x, z; w_n))$ by $a(w')$ when $w_c(x, z; w_n) \leq w'$. When

$$w' - \eta \leq w_c(X, Z; w_n) \leq w'' + \eta,$$

the error for each (x, z) is bounded by $\epsilon/3$; when the above inequality does not hold, the error for each (x, z) is bounded by M . Therefore,

$$|Ta(w) - Ta(w_n)| < |b_n - b| |a(w'') - a(w')| + \frac{2\epsilon}{3}. \quad (1.8.10)$$

Now consider the function $l : [0, \bar{w}] \rightarrow \mathbb{R}$ defined by $l(\tilde{w}) = \tilde{w} - w'$. Then the same calculation shows that

$$\begin{aligned} Tl(w) &= (w'' - w')b; \\ |Tl(w_n) - (w'' - w')b_n| &< \frac{\bar{w}\epsilon}{3M} + \eta. \end{aligned}$$

Therefore,

$$(w'' - w')|b - b_n| \leq |Tl(w) - Tl(w_n)| + \frac{\bar{w}\epsilon}{3M} + \eta. \quad (1.8.11)$$

Since l is a linear function, and $w_c(x, z; w_n)$ converges to $w_c(x, z; w)$ in the weak*-topology, $Tl(w_n) \rightarrow Tl(w)$. There exists an N_2 such that for $n \geq N_2$,

$$|a(w'') - a(w')| |Tl(w) - Tl(w_n)| < \frac{\epsilon}{6(w'' - w')}. \quad (1.8.12)$$

Combining Eqs. (1.8.8), (1.8.10), (1.8.11), and (1.8.12), we see that when $n \geq \max\{N_1, N_2\}$,

$$|Ta(w) - Ta(w_n)| < \frac{|a(w'') - a(w')|}{w'' - w'} \left[|Tl(w) - Tl(w_n)| + \frac{\bar{w}\epsilon}{3M} + \eta \right] + \frac{2\epsilon}{3} < \epsilon.$$

This shows that $Ta(w_n) \rightarrow Ta(w)$.

Proof of Lemma 3. The first-order condition Eq. (1.3.6) and the continuity of g_H/g_L and f_G/f_B implies that the support of $V'(w_c(X, Z; w))$ is an interval at every on-path history. Therefore, the support of $\{V'(w_t) : t > 0\}$ is an interval $[k_l, k_h]$. We show that the right end point must be $V'(0)$, and the same argument can be used to show that the left end point is $V'(\bar{w})$.

Lemma 5. $k_l < k_h$ unless the equilibrium is trivial: $w_0 = 0$ and the agent is protected by limited liability.

Proof. Suppose that $k_l = k_h$. Then Propositions 1 and 2 and the strict concavity of u implies that b_t is the same at all on-path histories, and therefore the agent's incentive is always zero, so the only thing players can do is to repeat the static Bayesian Nash equilibrium. \square

Suppose that $k_h < V'(0)$. Let $w_l = \inf\{w : V'(w) < k_h\}$. Then $w_l > 0$, and $w_l \notin W_I$. First suppose that $\mathbf{e}(w_l) = (0, 0)$. Then the program Eqs. (1.3.1)-(1.3.3) at w_l has a unique (x, z) -independent optimal choice of w_c , which will be denoted by $w_c(w_l)$. If $k_h > -1/u'(0)$, then $w_c(w_l) = \delta^{-1}w_l > w_l$, and $V'(\delta^{-1}w_l) = V'(w_l)$ by Eq. (1.3.6), but this contradicting the construction of w_l . If $k_h < -1/u'(0)$, then $V(w_c(w_l)) = \delta^{-1}V(w_l) > V(w_l)$, which implies that $w_c(w_l) < w_l$. However, one can then extend V linearly to the right of w_l unless $w_l = \bar{w}$. The linear extension contradicts the construction of w_l , and that $w_l = \bar{w}$ violates Assumption 3. Therefore, $\mathbf{e}(w_l) \neq (0, 0)$. By Assumption 4, there exists $\epsilon > 0$ so that $\mathbf{e}(w_l)$ is strictly optimal for $w \in [w_l, w_l + \epsilon)$. Choose ϵ small enough so that

$$\epsilon \leq \min\{\delta^{-1}(1 - \delta)c_G, \bar{w} - w_l\}; \text{ and define} \quad (1.8.13)$$

$$\tilde{M} = \max\left\{-\frac{\sup_x 1 - g_L(x)/g_H(x)}{\inf_x 1 - g_L(x)/g_H(x)}, -\frac{\sup_x g_H(x)/g_L(x) - 1}{\inf_x g_H(x)/g_L(x) - 1}\right\}; \quad (1.8.14)$$

$$M = \max\{\tilde{M}, \tilde{M}^{-1}\}. \quad (1.8.15)$$

By construction, $V'(w_l + \epsilon) < V'(w_l)$. Since V is differentiable, there exists a $w_{l2} \in (w_l, w_l + \epsilon)$ so that

$$\frac{k_h - V'(w_{l2})}{k_h - V'(w_l + \epsilon)} = \frac{1}{1 + M}. \quad (1.8.16)$$

Lemma 6. If $w \in [w_l, w_{l2}]$, then with positive probability, $V'(w_c(x, z; w)) > k_h$.

Proof. By construction $e(w_l)$ is strictly optimal at w . The agent's incentive constraint in an environment θ in which $e_\theta(w_l) = 1$ can be written as

$$\int w_c(x, z; w)[g(x, 1) - g(x, 0)]f(z, \theta)dx dz \geq \frac{(1 - \delta)c_\theta}{\delta}. \quad (1.8.17)$$

Suppose that $w_c(x, z; w) \in [w_l, w_l + \epsilon]$ for almost all (x, z) . Then

$$\int w_c(x, z; w)[g(x, 1) - g(x, 0)]f_\theta(z)dx dz \leq \int \epsilon[g(x, 1) - g(x, 0)]^+ f(z, \theta)dx dz < \epsilon.$$

However, ϵ is less than the right hand side of Eq. (1.8.17) by construction, a contradiction. Therefore, $w_c(X, Z; w)$ is out of the interval $[w_l, w_l + \epsilon]$ with positive probability.

Suppose that $w_c(X, Z; w) > w_l + \epsilon$ with positive probability. Since the likelihood ratio functions are continuous, Eq. (1.3.6) implies that $V'(w_c(X, Z; w))$ achieves its maximum and minimum with probability zero. Therefore, there exist non-empty intervals (x_1, x_2) and (z_1, z_2) in the support of X_t and Z_t , respectively, such that $V'(w_c(x, z; w)) < V'(w_l + \epsilon)$ for $x \in (x_1, x_2)$ and $z \in (z_1, z_2)$. WLOG assume that $g_H(x_1) - g_L(x_1)$ and $g_H(x_2) - g_L(x_2)$ have the same sign. Suppose that $g_H(x_1) > g_L(x_1)$. Then by Eq. (1.3.6),

$$\inf_x \frac{V'(w_c(x, z; w)) - V'(w)}{V'(w) - V'(w_c(x_1, z; w))} = \inf_x \frac{g(x, 0) - g(x, 1)}{g(x_1, 1) - g(x_1, 0)} \frac{\mu_G f(z, G)g(x_1, e_G(w)) + \mu_B f(z, B)g(x_1, e_B(w))}{\mu_G f(z, G)g(x, e_G(w)) + \mu_B f(z, B)g(x, e_B(w))} > \tilde{M}$$

by the monotonicity of the second fraction in z and the definition of \tilde{M} in Eq. (1.8.14).

Therefore, with positive probability,

$$\begin{aligned} V'(w_c(x, z; w)) &> V'(w) + \frac{V'(w) - V'(w_c(x_1, z; w))}{\tilde{M}} \\ &> V'(w) + \frac{1}{M}[V'(w) - V'(w_l + \epsilon)] \\ &\geq \frac{1 + M}{M}V'(w_l) - \frac{1}{M}V'(w_l + \epsilon) \\ &= k_h. \end{aligned}$$

The last step uses Eq. (1.8.16). The case where $g(x_1, 1) < g(x_1, 0)$ can be treated in the same way.

If $w_c(x, z; w) < w_l$ with positive probability, then $V'(w_c(x, z; w)) > V'(w_l) = k_h$ with positive probability by the bang-bang property in Proposition 5. \square

By construction of w_l , Player 1's expected payoff at on-path histories reaches $[w_l, w_{l2}]$ with positive probability. This implies that with positive probability the slope of V' at the agent's expected payoff in the next period is strictly bigger than k_h , a contradiction. Therefore, $k_h = V'(0)$. The same argument shows that $k_l = V'(\bar{w})$.

Now we know that $k_h = V'(0)$ in any optimal equilibrium. Define w_l , ϵ , and w_{l2} as before. The above lemma now shows that Player 1's expected payoff is zero at the beginning of the second period with positive probability if the equilibria starts at some $w \in [w_l, w_{l2}]$, and thus with positive probability (eventually) for equilibria that start anywhere on the frontier.

Proof of Theorem 1.

We first prove the desired result for optimal equilibria that start at $w \notin W_I$. Define the measure ψ to be the Lebesgue measure on $[0, \bar{w}] \setminus W_I$ if the agent is not protected by limited liability, and define ψ to be the unit mass at zero if the agent is protected by limited liability. By Assumption 3, the interior of $[0, \bar{w}] \setminus W_I$ contains a neighborhood of \bar{w} , so ψ is not trivial and its support has a non-empty interior when the agent is not protected by limited liability. Lemma 3 implies that from every initial state w_0 , the state eventually enters every set A with positive ψ -measure. Therefore, the Markov process is ψ -irreducible. It also has the Feller property by Lemma 2. Propositions 6.2.5 and 6.2.8 of MT implies that the Markov process is a T -chain if the agent is not protected by limited liability. When he is protected by limited liability, V is linear in a neighborhood of 0, so there exists an $\epsilon > 0$ such that $\inf_{w \in (0, \epsilon)} P(w, \{0\}) > 0$, which means that $(0, \epsilon)$ is a small set of the Markov process. By Propositions 6.2.5 and 6.2.8 of MT again, the Markov process is a T -chain in this case too. Lemma 3 implies that the state zero is always reachable. The state space is compact, so Theorem 18.3.2 of MT implies that the Markov process is positive Harris recurrent. Finally, the process is aperiodic since $P(0, \{0\}) > 0$ whether or not the agent is protected by limited liability. Therefore, Theorem 2 applies, and proves the convergence of $P^n(w, \cdot)$ to the unique invariance distribution of the Markov process for all $w \in [0, \bar{w}]$. This proves the theorem for the case where $w_0 \in [0, \bar{w}] \setminus W_I$.

When $w \in W_I$, $P(w, \cdot)$ may not describe the true distribution of the continuation payoff. We need a separate argument to show that the distribution of w_t converges to π . Each optimal equilibrium can be associated with a $T = \min\{t > 0 : w_t \notin W_I\}$. Then T is

a stopping time. Moreover, w_t stays in a single connected component of W_I for $t < T$, and $w_t \in W_I$ with probability zero for $t \geq T$. By what we have shown, for each $n > 0$, $\sup_{A \in \mathcal{B}([0, \bar{w}])} |\text{Prob}(w_t \in A | T = n) - \pi(A)| \rightarrow 0$ as $t \rightarrow \infty$. To prove the theorem, it remains to show that either $T < \infty$ with probability one, or $\text{Prob}(w_t \in A | T = \infty) = \pi(A)$ for all A .

Lemma 7. Let (w', w'') be a connected component of W_I . Then $(e_B, e_G) = (0, 0)$ is not optimal anywhere on (w', w'') unless the agent is protected by limited liability and $w' = 0$.

Proof. Suppose that $(e_B, e_G) = (0, 0)$ is optimal at some $w \in (w', w'')$. First consider the possibility that $V'(w) \geq 0$, which happens only if the agent is protected by limited liability. Then the up front payment is zero by Proposition 2, so $w_c(x, z; w) = \delta^{-1}w$, and thus $V(w) = \delta V(\delta^{-1}w)$. However, this means that $(0, 0)$, $(w, V(w))$, and $(\delta^{-1}w, V(\delta^{-1}w))$ are on the same line, which implies that $w' = 0$.

Now assume that $V'(w) < 0$. Let $[w_{0l}, w_{0h}]$ be the maximal connected set that contains w on which $e_B = e_G = 0$ is optimal everywhere. Then $V'(w_{0h}) < 0$. The program Eqs. (1.3.1)-(1.3.3) at w_{0h} has a unique (x, z) -independent optimal choice of w_c , which will be denoted $w_c(w_{0h})$. Since $V(w_{0h}) \leq \delta V(w_c(w_{0h}))$, so $w_c(w_{0h}) < w_{0h}$. The first-order condition Eq. (1.3.6) further implies that V is linear on $[w_c(w_{0h}), w_{0h}]$. However, one can then extend $\mathcal{F}_{(0,0)}V$ linearly to the right of w_{0h} unless $w_{0h} = \bar{w}$. That $w_{0h} = \bar{w}$ is ruled out by Assumption 3, and the linear extension contradicts the construction of w_{0h} . \square

Let (w', w'') be the connected component that contains the initial state w_0 . If $e_B = e_G = 0$ is optimal somewhere on (w', w'') , then $w' = 0$ and the agent is protected by limited liability. Hence, $b_t = 0$ for all t when $T = \infty$. Clearly, the unique invariance distribution of the Markov process is the unit mass at zero, so when $T < \infty$, $w_t \rightarrow 0$ in probability. Therefore, $w_t \rightarrow 0$ in probability unconditional on T . In what follows, assume that $e_B = e_G = 0$ is never optimal on (w', w'') .

The strict concavity of u implies that $b_t = b' \equiv (u')^{-1}(-1/V'(w'))$ as long as $w_t \in [w', w'']$. If $T \geq n$ with probability one, then the agent can guarantee himself an expected payoff of $(1 - \delta^n)b_t$ by not exerting any effort in any period, and receives payoff at most $(1 - \delta^n)b_t + \delta^n \bar{w}$. Therefore, his incentive of exerting effort is at most $\delta^n \bar{w}$, which

approaches zero as $n \rightarrow \infty$. By continuity, there exists $N > 0$ and $\epsilon > 0$ such that if $\text{Prob}(T \geq N) > 1 - \epsilon$, then the principal cannot induce the agent to exert effort even in the good environment. Therefore, for every optimal equilibrium starting on $[w', w'']$, $\text{Prob}(T \geq N) \leq 1 - \epsilon$. It follows by induction that $\text{Prob}(T \geq kN) \leq (1 - \epsilon)^k$, so $\text{Prob}(T = \infty) = 0$.

Proof of Proposition 6. Fix an open subset W of $[0, \bar{w}] \setminus W_I$. Then Lemma 3 implies that every optimal equilibrium enters W with positive probability. However, if the distribution of the initial state is π , the distribution of w_t is π for all t . Therefore, π must assign positive probability to W . The same argument shows that $\pi(\{0\}) > 0$ and $\pi(\{\bar{w}\}) > 0$. Finally, the proof of Theorem 1 implies that every optimal equilibrium enters $[0, \bar{w}] \setminus W_I$ and stays there with probability one. Therefore, the support of π is $[0, \bar{w}] \setminus W_I$.

Proof of Proposition 7. First assume that a stationary equilibrium exists. Then we have seen that $\bar{w} = \bar{s}$ and a stationary equilibrium exists for every $w \in [w_m, \bar{w}]$. Now consider the program Eqs. (1.5.5)-(1.5.6). Suppose that the value is strictly less than w_m . Then there also exists a stationary equilibrium at that value, contradicting the definition of w_m . Now let $\tilde{w}_m = \inf\{w : V'(w) = -1\}$. Suppose $\tilde{w}_m < w_m$. Then since $\tilde{w}_m + V(\tilde{w}_m) = \bar{s}$, the optimal equilibrium at \tilde{w}_m must induce effort $\mathbf{e} = (1, 1)$ from the agent in every period on the equilibrium path. In particular, $\mathbf{e}(\tilde{w}_m) = (1, 1)$. Consider the program Eqs. (1.3.1)-(1.3.3) and a number $\lambda \in (0, 1)$. Then $\lambda w_c(x, z; \tilde{w}_m) + (1 - \lambda)w_c(x, z; w_m)$ is feasible at $\lambda \tilde{w}_m + (1 - \lambda)w_m$, and it gives the principal payoff strictly higher than $\lambda V(\tilde{w}_m) + (1 - \lambda)V(w_m)$ unless V is linear on the convex hull of $\{w_c(x, z; \tilde{w}_m), w_c(x, z; w_m)\}$ for almost all (x, z) . However, this means that there exists a stationary equilibrium at \tilde{w}_m , contradicting the definition of \tilde{w}_m . Therefore, $V'(w) > -1$ for $w < w_m$.

Now assume that w_m solves Eqs. (1.5.5)-(1.5.6). Clearly, the solution of the linear program has the bang-bang property: $w_c(x, z) \in \{w_m, \bar{s}\}$ for almost all (x, z) . Let $A = \{(x, z) : w_c(x, z) = \bar{s}\}$. Consider the following automaton. There are two states, 1 and 2. In State 1, the principal makes up-front payment b_0 , and in State 2, the principal makes a payment of $(1 - \delta)^{-1}(\bar{s} - w_m) + b_0$. In both states, the agent exerts effort in both environments, and the state in the next period is 2 if and only if $(X, Z) \in A$. The game starts at state 1. By definition of A , the strategy profile satisfies the agent's IC

constraint Eq. (1.3.3). It will be an equilibrium if the principal's continuation payoff is nonnegative in both states. Let v_1 be her expected payoff at State 1. Then her payoff at State 2 is $v_1 - \bar{s} + w_m$. By Eq. (1.3.1),

$$v_1 = (1 - \delta)(y - b_0) + \delta[v_1 - \text{Prob}((X, Z) \in A | e_B = e_G = 1)(\bar{s} - w_m)].$$

By Eq. (1.5.5),

$$w_m = (1 - \delta)(u(b_0) - E[c_\Theta]) + \delta[w_m + \text{Prob}((X, Z) \in A | e_B = e_G = 1)(\bar{s} - w_m)].$$

Adding the above two equations together leads to $v_1 = y - E[c_\Theta] - w_m + u_0 = \bar{s} - w_m$. Therefore, the principal's payoff in State 1 is $\bar{s} - w_m > 0$, and her payoff in State 2 is 0. The proposed automaton is indeed an equilibrium.

Notice that ω is a convex function and is thus continuous. Also, $\omega(0) > 0$ as long as $\delta \int \bar{s}[g(x, 1) - g(x, 0)]^+ dx > (1 - \delta)c_B$. Therefore, to show that ω has a fixed point, it suffices to show that $\omega(w) < w$ for some $w > 0$. Let

$$w_B = \bar{s} - \frac{(1 - \delta)c_B}{\delta \int [g(x, 1) - g(x, 0)]^+ dx}.$$

Then when $w = w_B$, there exists a unique w_c that satisfies Eq. (1.5.6): $w_c(x, z) = w_B$ for all low x and $w_c(x, z) = \bar{s}$ for all high x . Now

$$\omega(w_B) = (1 - \delta)(u(b_0) - E[c_\Theta]) + \delta w_B + \delta[\bar{s} - w_B] \text{Prob}(g(X, 1) > g(X, 0) | e_B = e_G = 1).$$

Therefore, $\omega(w_B) < w_B$ if

$$w_B > (u(b_0) - E[c_\Theta]) + \frac{\delta c_B \text{Prob}(g(X, 1) > g(X, 0) | e_B = e_G = 1)}{\int [g(x, 1) - g(x, 0)]^+ dx}.$$

Notice that the right hand side is independent of \bar{s} , while the left hand side is linearly increasing in \bar{s} . Therefore, the $w_B > \omega(w_B)$ for sufficiently big \bar{s} .

Proof of Proposition 8. Define a binary relation \sim on $[0, \bar{w}]$ by setting $w' \sim w''$ if and only if $w' = w''$ or $w', w'' \in \{0\} \cup [w_m, \bar{w}]$. Let $S = [0, \bar{w}] / \sim$ with the topology of the circle. Denote the image of 0 under the quotient map by s_0 . Then the Markov transition probability $\tilde{P}(s, \cdot)$ is well defined in the new state space S for $s \neq s_0$. Define $\tilde{P}(s_0, \cdot)$ to be the unit mass at s_0 . Then it is easy to verify that Lemmas 2 and 3 hold

for the Markov process on S as long as the initial state in Lemma 3 is not s_0 . Notice that a continuous function a on $[0, \bar{w}]$ induces a continuous function on S if and only if $a(0) = a(w)$ for all $w \in [w_m, \bar{w}]$. It is straight forward to show that the Markov process is ψ -irreducible when ψ is the unit mass at s_0 .

Since V is linear in a neighborhood of 0, and the proof of Lemma 3 implies that the probability that $w_c(X, Z; w) \geq w_m$ is bounded from below for w in a neighborhood of w_m . Therefore, a neighborhood of s_0 in S is a small set of the Markov transition probability \tilde{P} . Therefore, the Markov process is a T -chain. The state s_0 is reachable, so the process is positive Harris recurrent. Finally, since $\tilde{P}(s_0, D) = 0$ for all D that does not include s_0 , the process is aperiodic. Therefore, the aperiodic ergodic theorem applies. Clearly, the unique invariant distribution is the unit mass at s_0 .

Proof of Proposition 10. Fix a $w \in [0, \bar{w}] \setminus W_I$. Then the right hand side of Eq. (1.6.4) is an analytic function of z . When the analytic function is a constant, $w_c(z; w)$ lies in a connected component (w', w'') of W_I for almost all z . Since $w \notin W_I$, there exists a sequence $w_n \in [0, \bar{w}] \setminus (W_I \cup \{w\})$ that converges to w . WLOG, assume that the $(w - w_n)$'s all have the same sign. When $w_n < w$ for all n , V cannot be linearly extended to the left of w , so

$$\begin{aligned} w &= \min_{\bar{w}_c} (1 - \delta)(h - E[c_{\Theta} e_{\Theta}]) + \delta E[\bar{w}_c(Z)|e] \\ \text{s.t. } \delta \int w_c(z)[f_1(z, \theta) - f_0(z, \theta)] dz &= (1 - \delta)c'_{\theta}(e_{\theta}). \end{aligned}$$

This implies that $w_c(z) \in \{w', w''\}$ for almost all z . When $w_n > w$ for all n , w is the solution of the above program with \min replaced by \max , and we will have $w_c(z) \in \{w', w''\}$ for almost all z . Therefore, $w_c(z)$ is at the boundary of W_I for almost all z when the right hand side of Eq. (1.6.4) is a constant. When the right hand side of Eq. (1.6.4) is not a constant, $\{z : V'(w_c(z; w)) = k\}$ is a discrete set for every k since the zeros of an analytic function are isolated. On the other hand, $V'(W_I)$ is a discrete set, so $\{z : w_c(z; w) \in W_I\}$ is a countable set and thus $w_c(Z; w) \in W_I$ with probability zero.

As before, define the transition probability by setting $P(w, W) = \text{Prob}(w_c(Z) \in W|e(w))$ for $w \notin W_I$, and using convex combination for $w \in W_I$. In the proof of Lemma 2, we need to focus on the case where $e(w_n)$ converges. Then a $(w_c(\cdot; w_n), V(w_c(\cdot; w_n)))$ that converges in the weak*-topology on $L^\infty(\Omega)$ gives an optimal choice of w_c in the program

Eqs. (1.6.1)-(1.6.3) at w . The rest of the proof of Lemma 2 can be used without changing a word. Since the right hand side of Eq. (1.6.4) is continuous and Ω is connected, the support of $w_c(Z; w)$ is connected for all w . The proof of Lemma 3 is based on the boundedness of likelihood ratio from zero and infinity and the strict concavity of u . These are assumed in this setting too, so Lemma 3 also holds.

Therefore, the Markov process is has the Feller property and is ψ -irreducible where ψ is the Lebesgue measure on $[0, \bar{w}] \setminus W_I$. The support of ψ contains a neighborhood of \bar{w} by Assumption 3, so the Markov process is a T -chain. The state zero is reachable and the state space is compact, so the process is positive Harris recurrent. Since $P(0, \{0\}) > 0$, the process is aperiodic. Now the proposition holds when $w_0 \notin W_I$ by Theorem 2. The case where $w_0 \in W_I$ can be handled the same way as in the proof of Theorem 1.

1.8.2 Appendix B: long-run dynamics with commitment

In this appendix, we study the long-run dynamics when the principal can commit to a long-term contract. For simplicity, we focus on the case where the agent is not protected by limited liability. Without the fluctuating environment Θ_t , this will be the discrete time analogue of Sannikov (2008) without limited liability. The analysis is not much different from the no-commitment case: it only removes the constraint that $V \geq 0$ so the agent's expected utility w can potentially be arbitrarily large. However, we will see that the long-run dynamics will be very different. In particular, the state w_t will eventually be absorbed to a counter-intuitive "retirement" regime.

To begin with, we first estimate an upper bound of the frontier V by ignoring the agent's incentive-compatibility constraint for a moment. The only constraint is feasibility then.

We have $V(w) \leq V_f(w)$, where

$$\begin{aligned} V_f(w) &= \max_{q_B, q_G, h} -u^{-1}(h) + \sum_{\theta} \mu_{\theta} q_{\theta} y; \\ s.t. & \quad h - \sum_{\theta} \mu_{\theta} q_{\theta} c_{\theta} = w. \end{aligned}$$

In this program, q_θ is the probability that the agent exerts effort in Environment θ . Eliminating h , we see that q_B and q_G maximizes

$$-u^{-1}\left(w + \sum_{\theta} \mu_{\theta} q_{\theta} c_{\theta}\right) + \sum_{\theta} \mu_{\theta} q_{\theta} y.$$

We make the following assumption:

Assumption 6. $\lim_{b \rightarrow \infty} u'(b) < c_G/y$.

Both the commonly used exponential utility and power function utility satisfies this condition. Under this assumption, for sufficiently large w , it is optimal to choose $q_B = q_G = 0$, so $V_f(w) = -u^{-1}(w)$. Since there is no incentive-compatibility constraint to worry about for these w , the payoff pair $(w, u^{-1}(w))$ can be implemented by a long-term contract. Therefore, $V(w) = -u^{-1}(w)$ for sufficiently large w . Let $w_r = \inf\{w : V(w) = -u^{-1}(w)\}$.

The solution to the program Eqs. (1.3.1)-(1.3.3) is still characterized by the first-order condition Eq. (1.3.6). Define $W_I = \{w \in (0, w_r) : V \text{ is linear on a neighborhood of } w\}$. Assume that Assumptions 2, 4, and 5 hold, and assume in addition that 0 is not in the closure of W_I . Then Proposition 5 holds, and one can define the Markov transition probability P as in Section 5.2.

Proposition 11. In every optimal equilibrium, w_t converges to a random variable w_∞ almost surely, and $w_\infty \geq w_r$ almost surely.

Proof. WLOG assume that $w_0 \leq w_r$. Since zero is the only possible corner solution of Eq. (1.3.6), $V'(w_t)$ is a super-martingale. Moreover, since $g(x, 1)/g(x, 0)$ is bounded away from zero and infinity, $w_c(x, z; w)$ is bounded as long as the Lagrange multipliers κ_B and κ_G are bounded. Clearly, there exists a $K > 0$ such that if $|\kappa_B| > K$ or $|\kappa_G| > K$, then for all $w \in [0, w_r]$, $w_c(x, z; w)$ cannot satisfy either IC constraint with equality. Therefore, there exists $w_M < \infty$ such that $w_c(x, z; w) \leq w_M$ for all $w \in [0, w_r]$. Moreover, $w_c(x, z; w) = w$ for $w > w_r$. Therefore, $w_t \leq w_M$ almost surely. By Doob's martingale convergence theorem, $V'(w_t)$ converges a.s. and in L^1 to some random variable k_∞ .

Suppose that there exists a $k \in (V'(w_r), V'(0))$ such that $k_\infty = k$ with positive probability $\xi > 0$. Let $[w', w''] = (V')^{-1}(k)$. Proposition 5 and the argument in the last

paragraph of the proof of Theorem 1 applies, and implies that $w_t \in (w', w'')$ with probability zero for all t . Lemma 3 also applies when the initial state is in $(0, w_r)$, and implies that the state reaches $\{0\}$ with positive probability when the process starts at w' or w'' . Define the operator T as in Eq. (1.5.1). Then there exists $n > 0$ and $\epsilon > 0$ such that $T^n V'(w') - V'(w') < -2\epsilon$ and $T^n V'(w'') - V'(w'') < -2\epsilon$. Since V' is continuous and the Markov process P has the Feller property, $T^n V'$ is also continuous, and there exists $\eta > 0$ such that $T^n V'(w) - V'(w) < -\epsilon$ for $w \in (w' - \eta, w') \cup [w'', w'' + \eta)$. Since $V'(w_t)$ converges to k_∞ a.s., there exists $t_k > 0$ such that $\text{Prob}(w_t \in (w' - \eta, w') \cup [w'', w'' + \eta)) > \xi/2$ for all $t > t_k$. Therefore,

$$E[V'(w_{t+n})] = E[T^n V'(w_t)] < E[V'(w_t)] - \frac{1}{2}\xi\epsilon,$$

for all $t > t_k$, where we have used the fact that $T^n V'(w_t) = E[V'(w_{t+n})|w_t]$. However, $E[V'(w_t)]$ converges to $E[k_\infty]$, a contradiction. Therefore, $k_\infty = k$ with zero probability for every $k \in (V'(w_r), V'(0))$. In particular, $k_\infty \in V'(W_I)$ with zero probability. Hence, $w_\infty = (V')^{-1}(k_\infty)$ is well-defined up to events with probability zero, and w_t converges to w_∞ a.s..

Define $q : [0, \infty) \rightarrow [0, w_r]$ by $q(w) = w$ for $w \in [0, w_r]$ and $q(w) = w_r$ for $w > w_r$. Define the Markov transition probability \tilde{P} on $[0, w_r]$ by

$$\tilde{P}(w, W) = \begin{cases} P(w, W), & \text{if } w_r \notin W; \\ P(w, W \cup [w_r, \infty)), & \text{if } w_r \in W. \end{cases} \quad (1.8.18)$$

Let \tilde{T} be the T -operator defined in Eq. (1.5.1) for \tilde{P} . Each continuous function a on $[0, w_r]$ can be extended to $[0, \infty)$ by setting $a(w) = a(w_r)$ for $w > w_r$, and then

$$\tilde{T}a(w) = \int a(w')\tilde{P}(w, dw') = \int a(w')P(w, dw')$$

by construction. Lemma 2 still applies and says that P has the Feller property. The above equation implies that \tilde{P} also has the Feller property, and \tilde{T} maps a continuous function to a continuous function.

Let π be the probability distribution of w_∞ , and

$$\tilde{\pi}(W) = \int \tilde{P}(w, W)d\pi(w).$$

Suppose that $\pi \neq \tilde{\pi}$. Since π and $\tilde{\pi}$ are both finite measures on the compact metric space $[0, w_r]$, they are both regular by Theorem 7.1.4 of Dudley (2002). By Theorem 7.4.1 of Dudley (2002), there exists a continuous function a_0 on $[0, w_r]$ such that $\int a_0(w)d\pi(w) \neq \int a_0(w)d\tilde{\pi}(w)$. Since a_0 is continuous and bounded,

$$\int a_0(w)d\pi(w) = E[a_0(q(w_\infty))] = \lim_{n \rightarrow \infty} E[a_0(q(w_{t_n}))]; \quad (1.8.19)$$

$$\int a_0(w)d\tilde{\pi}(w) = \int \tilde{T}a_0(w)d\pi(w) = \lim_{n \rightarrow \infty} E[\tilde{T}a_0(q(w_{t_n}))]. \quad (1.8.20)$$

However, for each t ,

$$E[\tilde{T}a_0(q(w_t))] = E\left[\int a_0(w')\tilde{P}(q(w_t), dw')\right] = E[a_0(q(w_{t+1}))], \quad (1.8.21)$$

where we used the fact that $\tilde{P}(q(w_t), \cdot) = P(w_t, \cdot)$, which follows from the fact that $w_c(x, z; w) = w$ for $w \geq w_r$. Combining Eqs. (1.8.19)-(1.8.21) shows that $\int a_0(w)d\pi(w) = \int a_0(w)d\tilde{\pi}(w)$, a contradiction. Therefore, π is an invariant distribution of the Markov process on $[0, w_r]$ with transition probability \tilde{P} .

It is easy to check that Lemma 3 holds for \tilde{P} , which implies that every invariant probability measure of \tilde{P} assigns positive probability to $\{w_r\}$. (See also the discussion of the case without commitment.) By Theorem 1.7 of Hairer (2010), \tilde{P} has a unique probability measure, the unit mass at w_r . Therefore, $w_\infty \geq w_r$ with probability one. \square

The proof is an application of Doob's martingale convergence theorem. Notice that w_∞ depends on the initial state w_0 in general, but for every initial state, $\text{Prob}(w_t \geq w_r)$ converges to one as $t \rightarrow \infty$.

Therefore, as long as Assumption 6 holds, the relationship almost surely eventually enters a regime where the agent does not exert any effort and the principal makes a constant payment every period. This is exactly the "retirement" regime in Sannikov (2008). Notice that $V(w) < 0$ for all $w \geq w_r$, so "retirement" can never happen when the principal cannot commit to a long-term contract. However, the principal's commitment power makes all information about the invariant distribution of the Markov process in 4.2 lost in the long run. Therefore, the model with commitment is not a good approximation of the model without commitment in terms of long-run dynamics.

In Sannikov's model, both zero and the retirement regime are absorbing. However, when the agent is not protected by limited liability, the state zero is not absorbing. In fact, when the agent is protected by limited liability, the martingale convergence argument used in the proof of the proposition implies that w_t converges to some random variable w_∞ , and $w_\infty \in \{0\} \cup [w_r, \infty)$ with probability one. The argument in the proof of proposition 8 shows that $w_\infty = 0$ with positive probability and $w_\infty \geq w_r$ with positive probability, if $w_0 \in (0, w_r)$.

Chapter 2

Relationships in a Market: Contract Enforcement and Asset Re-allocation

}

2.1 Introduction

In a principal-agent relationship, when the agent's performance is not verifiable by a third party, a formal incentive contract cannot be written. If his performance is observable by both the principal and the agent and they interact repeatedly, the principal can still provide incentive to the agent through a "relational contract". In general, the lack of formal contract reduces the surplus that can be created by the relationship. However, the reduction in pair productivity can benefit the whole economy if it facilitates the transfer of assets from inefficient managers to efficient managers by making it harder for less efficient managers to survive.

In order to formally explore this possibility, this paper studies a market of relationships. Specifically, each relationship consists of a principal, an agent and an asset, and assets can be traded in the market. The economy consists of a continuum of principals, agents, and assets, and the productivity of a relationship depends both on the agent's effort and on the matching quality between the asset and the relationship which is unknown until the principal owns the asset. Under some conditions, the productivity of each

relationship is constant over time, and the growth of the aggregate output of the whole economy comes from the trading of assets and the increase in the average matching quality in all relationships associated with the trading.

The surplus in every relationship depends on the contracting environment. As mentioned above, the lack of formal incentive contracts reduces the surplus in every relationship. However, if this reduction is particularly big for relationships in which the matching qualities between the principals and the assets are low, more principals sell their assets in the economy without formal incentive contract. Consequently, the limit distribution of matching qualities is better in this economy. It is therefore interesting to ask whether this effect dominates the reduction of output in the productive relationships and leads to higher output and higher welfare in the economy without formal contract.

There are many reasons why the initial allocation of assets may be sub-optimal. For example, there may be friction in the credit market so that the initial asset owners are those who can afford assets and not those who are good at managing them. This paper does not explicitly model these frictions. Instead, it assumes that the asset market is frictionless except for a deadweight loss associated with every trade, and explores the dynamics of asset re-allocation in both the environment with formal contracts (the strong contracting environment) and the environment without formal contracts (the weak contracting environment).

Specifically, there are continuums of entrepreneurs (principals), workers (agents) and assets. Each asset is initially owned by an entrepreneur. In each period, the entrepreneur can either sell her asset on an asset market or keep her asset and hire a worker. The assets are scarce resources so that current asset owners have all the bargaining power in the asset market and the entrepreneur who owns an asset has all the bargaining power in the contract with the worker. After a transaction of asset takes place, the asset remains unproductive for the current period, and the new owner learns the matching quality between herself and the asset. When the entrepreneur hires a worker, the worker chooses an effort that is not observed by the entrepreneur, but the effort generates a binary output which is publicly observable. An important assumption is that the worker's output and the entrepreneur's payment are observed by every agent in the economy. Therefore, if an entrepreneur reneges her payment in the weak contracting environment, she can be punished even if she hires a new worker.

There are three main sets of results. The first set is concerned with whether the asset market is stationary. The market is called stationary if the asset price offered by each asset owner is constant over time. In the strong contracting environment, the asset market is always stationary. In the weak contracting environment, the asset price is a fixed point of an increasing function, and the asset market is stationary if that increasing function has a unique fixed point. Intuitively, since the current asset owner has all the bargaining power, she will offer the price equal to the value of the asset to a new owner. Since the potential buyer does not know her matching quality with the asset, there is no private information in the asset market. However, in the weak contracting environment, the value of the asset may depend on the future asset prices in a way that leads to a non-constant asset price.

The second set of results characterizes optimal relationships in a stationary asset market. In such an asset market, a new asset owner either sells her asset immediately or keeps the asset forever. Moreover, when she keeps her asset forever, the surplus created from the asset is constant over time. The intuition is simple: since the asset price will not change in a stationary asset market, there is no point of keeping an asset if the owner finds that her matching quality with the asset is low. It is obvious that the relationship in the strong contracting environment produces constant surplus, while the same result in the weak contracting environment follows from Levin (2003).

Finally, in a stationary asset market, the aggregate output of the economy converges to a limit in the long run, and the distance between the output at time t and the limit shrinks exponentially over time. As mentioned above, the growth of the aggregate output is driven by the re-allocation of assets. Since the asset owners whose matching qualities with their assets are above a threshold keep their assets forever, eventually all assets fall into hands with high enough matching qualities and the aggregate output reaches its long-run limit. This long-run limit can be characterized analytically, so it can be shown when this long-run limit is higher in the weak contracting environment.

This paper is related to several literatures. The literature of relational contracts considers incentives in the weak contracting environment. Examples include Macleod and Malcomson (1989), Baker, Gibbons, and Murphy (1994, 2002) and Levin (2003). Most of these papers consider a single relationship. By considering an asset market, the current paper shows that factors that reduce productivity of every relationship may benefit

the economy as a whole.

There are also papers that relate the performance of a single relationship to market forces, an early example of which is Shapiro and Stiglitz (1984), and there are also recent papers like McAdams (2011) and Halac (2012). Kranton (1996) and Eeckhout (2006) also study relationships with re-matching. However, none of these papers have a market of relationships or assets and pose the problem in the context of economic growth.

In macroeconomics, Jovanovic (1979 and 1982) are early papers studying productivity shocks that motivate the consideration of matching qualities between entrepreneurs and assets in this paper. Bloom (2009) shows through a numerical simulation the contribution of this re-allocation process to growth of economies. Hsieh and Klenow (2009) use establishment-level data in India and China to show the potential of growth in those countries through re-allocation. Song, Storesletten, and Zilibotti (2011) argue that the initial inefficient allocation of resources is due to the inefficiencies in the credit market, so the resources initially belong to firms who have more cash than those who are more productive.

The remainder of the paper is organized as follows. Section 2 outlines the assumptions of the model. Section 3 analyses the Pareto optimal formal and informal incentive contracts between one employer and one employee, taking the market conditions as exogenous. Section 4 analyses the full market equilibrium and characterizes the growth path of the economy. Section 5 describes an extension that allows for a continuous measure of enforceability of incentive contracts in the spirit of Baker (1992) and Baker (2002). Section 6 concludes. The proofs that are not included in the main text are collected in the appendix.

2.2 The Model

In a discrete time economy, there is a continuum of assets of mass one, a continuum of entrepreneurs of mass $N > 1$, and a continuum of workers of mass $N_w \in (1, N)$. Each entrepreneur can own at most one unit of asset, and hire at most one worker in each

period of time¹. Each entrepreneur has an matching quality a with each asset, which is independent across entrepreneurs and assets and drawn from a distribution at the beginning of the game, whose CDF is F with support $[\underline{a}, \bar{a}]$, where $\underline{a} > 0$. Allowing each entrepreneur to own more than one unit of assets naturally raises the question of whether productivity of different assets owned by the same entrepreneur are the same, and thus distinguishes the interpretations of a as the entrepreneur's innate ability or as the quality of the matching between the entrepreneur and each asset. This will be left for future research.

At the beginning of the first Period, Period 0, each unit of asset is owned by an entrepreneur, and the distribution of the initial matching qualities is given by the cumulative distribution function F_0 , which is taken as exogenous and may not be the same as F . This assumption captures the idea that the initial allocation of assets is inefficient. This inefficiency may be caused by friction in the credit market as described in Song, Storesletten and Zilibotti (2011), but is not explicitly modeled here, and after the initial allocation of assets there is no more friction in the credit market in the sense that every entrepreneur can afford an asset. I also assume that each entrepreneur observes her matching quality with an asset only after she owns the asset. This assumption prevents the allocation of assets becoming efficient after a single period during which only entrepreneurs with the highest abilities bid to buy the assets previously owned by entrepreneurs with lower abilities. Alternatively, a can be interpreted as the quality of the matching between an entrepreneur and an asset, and thus unknown prior to the formation of the match. In each period of time, the following events happen:

1. The matching quality between each asset owner and her asset is observed by everyone in the economy.
2. Each entrepreneur who owns an asset decides whether to sell her asset.

¹Allowing each entrepreneur to hire more than one worker makes the characterization of optimal incentive contracts more difficult. When the output of the firm is the sum of contribution from all its employees, the optimal formal incentive contract for each employee is the set of optimal incentive contract based on his own contribution, but the optimal informal incentive contract has a more complicated structure. The friction introduced by the informal incentive contracts remains the same, though somewhat alleviated due to the smaller fluctuation in total output, and vanishes when the number of employees in a firm approaches infinity. When the output of the firm depends on the contribution from each employee in a more general way, the problem of team incentive arises.

-
3. Each entrepreneur who decides to sell her asset is randomly matched with an entrepreneur who does not own an asset, and makes a take-it-or-leave-it offer. If it is accepted, the buyer gets the asset.
 4. Each entrepreneur who decides not to sell her asset offers a wage w and incentive payment schedule $m(\cdot)$ to her current employee (in case she employed someone in the previous period) or a randomly matched worker. Here m is a function of the worker's performance z to be defined later.
 5. For each matched pair of entrepreneur and worker, the worker chooses to accept or reject the wage offer. If the wage offer is rejected, both the entrepreneur receives zero payoff in the current period. Each unemployed worker receives an unemployment benefit of $b \geq 0$, no matter whether he is offered a contract.
 6. If a wage offer is accepted, the wage w is paid, and the worker chooses a level of effort $e \in [0, 1]$ at a private cost $c(e)$. We assume that c is a strictly increasing, twice continuously differentiable function. We also assume that $c(0) = 0$, $c'(0) \geq 0$, and $c''(e) > 0$ for all $e \in [0, 1]$.
 7. Within each pair of matched entrepreneur and worker, an outcome $z \in \{0, 1\}$ is realized and observed by every agent in the economy, where $\Pr(z = 1|e) = e$ where e is the effort chosen by the worker; the output is $y = az$ where a is the entrepreneur's matching quality with the asset, and makes an incentive payment.

Throughout the analysis, we assume that $c'(1) > \bar{a}$, which would imply that it is never optimal to choose effort $e = 1$ in equilibrium. All agents in the economy are risk neutral, and discount future payoffs by δ . In each period, an entrepreneur's payoff is p if she sells her asset at price p , $-p$ if she buys an asset at price p , zero if she does not own an asset, or has a contract offer turned down by a worker, and $(az - w - m)$ if she has a contract accepted by a worker and pays wage w and bonus m . In each period, a worker's payoff is b if he is not matched with an entrepreneur or rejects a contract offer, or $w + m - c(e)$ if he accepts a contract offer and exerts effort e .

Each worker's effort is observed by himself, and the matching quality a , the outcome z and monetary payment w and m that are offered and realized associated with each current owner of the asset is publicly observed. On the other hand, once an asset transaction takes place, information about its previous owner is forgotten. This means

that the resale value of an asset does not depend on actions taken by the current owner. Even if the current owner of an asset has deviated from the equilibrium path, the next owner can treat it as a new asset. This is similar to the "fresh start" assumption of McAdams (2011). The solution concept is public perfect equilibrium in pure strategy with one restriction, as stated in Assumption 1 later.

It follows from the timing that once matched, an entrepreneur and a worker will work together until the entrepreneur sells her asset or the worker rejects a contract offer, and until then they do not directly interact with any other agents. Therefore, it is possible to analyze the interaction between an entrepreneur and a worker, taking the market price of assets and the worker's payoff from being unemployed as given. I call this interaction the *partnership game* between the entrepreneur and the worker, and will study it in detail in the next section.

We distinguish between the strong contracting environment and the weak contracting environment. In the strong contracting environment, the entrepreneur has to make the promised incentive payment $m(z)$ upon the realization of z . In the weak contracting environment, the entrepreneur can choose whether to make the promised incentive payment. It is assumed that the entrepreneur has all the bargaining power in the negotiation of labor contracts to capture the idea that the labor force has a bigger mass than assets. As a result, all the surplus in the employment relationship accrues to the entrepreneur in the strong contracting environment, as will be shown in the next section. However, the same result is not necessarily true in the weak contracting environment, due to the multiplicity of equilibria in repeated games. To facilitate the comparison between the two contracting environments, we make the following assumption:

Assumption 7. In each employment relationship, the entrepreneur and the worker choose strategies on the Pareto frontier, unless at least one of the two parties have reneged in the current partnership, at which point they play the static Nash equilibrium if no one quits.

This assumption is similar to Assumption 2 in Halac (2013). It implies that an entrepreneur cannot be punished for choosing one equilibrium of the partnership game over another. However, she focuses on a single partnership, and I also study the formation of new partnership through new matching or transaction of assets.

To capture the idea that the entrepreneurs have a bigger mass than assets, it is assumed that in an asset transaction the seller has all the bargaining power. In fact, when the mass of entrepreneurs, N , is sufficiently big, and all entrepreneurs are *ex ante* the same, the chance that each entrepreneur will get an asset is slim. Therefore, assuming that the seller captures all the surplus from an asset transaction seems reasonable.

In this paper, we focus on the asset market, has a very simplistic model of the labor market, and completely leave out the product market of the economy. In particular, we will not discuss how entrepreneurs with assets compete against each other in the product market. As a result, an important aspect that is missing from the model is the effect of the distribution of asset owners' abilities on the price in the product market and thus the profit of all asset owners. In fact, as we will show in Section 4, the distribution of asset owners' abilities improves over time in the sense of first order stochastic dominance, and that may cause a drop in the price of the product market and reduce each asset owner's profit, but this is left for future research. Another consequence of competition in the product market is that the consumer surplus will be non-zero, and thus in general, the social welfare of the economy is not equal to the total monetary payoffs of entrepreneurs and workers. Due to this consideration, we will measure the performance of the whole economy by its aggregate output, the sum of az over all asset owners, instead of the total monetary payoffs of entrepreneurs and workers.

2.3 The Partnership Game and Determination of Asset Price

This section characterizes the equilibriums in both contracting environments. In particular, we will focus on whether the asset price is constant over time, and discuss properties of such an equilibrium. Section 3.1 studies the strong contracting environment, where the equilibrium is unique and the asset price is always constant over time, and Section 3.2 studies the weak contracting environment and gives conditions under which an equilibrium with constant asset price exists and is unique.

In order to focus on the partnership game between an entrepreneur and a worker, We first introduce a lemma stating that for the entrepreneur in a partnership, the resale

price of assets and the worker's outside option can be taken as parameters that are independent of her own strategy.

Lemma 8. Fix a contracting environment. In each period, all sellers of assets offer the same price. All entrepreneurs who does not own assets at the beginning receive total payoff of zero. In particular, the seller receives the whole surplus in an asset transaction. Similarly, all workers receive payoff $b/(1 - \delta)$, the payoff from being unemployment.

This lemma applies to both contracting environments, and fleshes out the implications of the assumptions about bargaining powers. Since the sellers have all the bargaining power in asset transactions, the price of assets is independent of the seller's value if the transaction breaks down, and is equal to the buyer's continuation value. Similarly, since the entrepreneur has all the bargaining power in wage negotiation, the worker cannot receive more than his outside option in expectation in an employment relationship. However, the lemma does not pin down the asset price in the market, nor does it imply that the price is constant over time. The remaining of this section is devoted to the determination of asset price and the value of an asset for entrepreneurs with different abilities.

2.3.1 The strong contracting environment

Let p_t be the equilibrium market price of assets at time t , and $V_t^{(s)}(a)$ the value of assets at time t in a matching of quality a . Then the asset owner chooses the contract (w, m) and recommended level of effort e to solve the following program:

$$\begin{aligned} V_t^{(s)}(a) &= \max_{w, m, e} \{p_t, [a - m_1]e - m_0(1 - e) - w + \delta V_{t+1}^{(s)}(a)\} \\ \text{s.t. } & e \in \operatorname{argmax}_{\bar{e}} [m_1 - m_0]\bar{e} - c(\bar{e}); \\ & w + m_1e + m_0(1 - e) - c(e) \geq b. \end{aligned}$$

Here m_1 is the bonus payment when $z = 1$ is realized, and m_0 is the bonus payment when $z = 0$ is realized. The first constraint is the worker's incentive compatibility constraint, and the second constraint is the worker's participation constraint. Using the transformation $\tilde{w} = w + m_0$ and $\tilde{m}_1 = m_1 - m_0$, one can set $m_0 = 0$ without loss of generality. Solving for w from the constraints, one can rewrite the entrepreneur's objective function as $\max\{p_t, a e - c(e) - b + \delta V_{t+1}^{(s)}(a)\}$. Therefore, the optimal effort

choice is $\min\{(c')^{-1}(a), 1\}$, which can be implemented by choosing $m_1 = a$. Let

$$g(a) = \max_{e \in [0,1]} ae - c(e). \quad (2.3.1)$$

Then

$$V_t^{(s)}(a) = \max\{p_t, g(a) - b + \delta V_{t+1}^{(s)}(a)\}.$$

By Lemma 1, the price of assets p_t is equal to the expected value of an asset, so $p_t = \delta \int V_{t+1}^{(s)}(\bar{a}) dF(\bar{a})$. Define an operator \mathcal{T} on $L^1(F)$, the integrable functions on $[\underline{a}, \bar{a}]$ with respect the probability measure F as follows:

$$\mathcal{T}(V) = \max\{\delta \int V(\bar{a}) dF(\bar{a}), g - b + \delta V\}. \quad (2.3.2)$$

Then $V_t^{(s)} = \mathcal{T}(V_{t+1}^{(s)})$.

Theorem 3. There exists a unique equilibrium in the strong contracting environment.

The equilibrium has the following properties:

a) the asset price p is constant over time, is independent of the initial distribution of asset owners' abilities and satisfies the following equation:

$$p = \delta \int \max\{(1 - \delta)^{-1}[g(a) - b], p\} dF(a); \quad (2.3.3)$$

b) the value of an asset in a matching with quality a is given by $\max\{p, (1 - \delta)^{-1}(g(a) - b)\}$; and

c) there exists $a_c^{(s)} \in [\underline{a}, \bar{a}]$ such that after receiving an asset, an entrepreneur sells the asset immediately if her matching quality with the asset a is lower than some threshold $a_c^{(s)}$, and never sells the asset otherwise. In fact, $a_c^{(s)}$ is given by the following equation:

$$\frac{1}{1 - \delta} [g(a_c^{(s)}) - b] = p. \quad (2.3.4)$$

Proof. It is easy to see that $\|\mathcal{T}(V') - \mathcal{T}(V'')\| \leq \delta \|V' - V''\|$ for $V', V'' \in L^1(F)$, so \mathcal{T} is a contraction on $L^1(F)$ and has a unique fixed point V_f in $L^1(F)$. Let $p = \delta \int V_f(a) dF(a)$ and $a_c^{(s)}$ be given by Eq. (2.3.4). Then the strategy profile described in the theorem is an equilibrium.

Suppose a strategy profile is an equilibrium. Then it is still an equilibrium for a different distribution of the initial asset owners' abilities, and neither p nor the value function

depends changes. Therefore, the set of equilibrium value functions is independent of the distribution of the initial asset owners' abilities. Let \mathcal{V} be the set of $V_1^{(s)}$'s that can arrive in equilibrium. Then $\mathcal{V} \subset L^1(F)$, and moreover $\|V\| \leq (1 - \delta)^{-1}g(\bar{a})$ for every $V \in \mathcal{V}$, since $V_1(a)^{(s)} \in [0, (1 - \delta)^{-1}g(\bar{a})]$ for all a and in every equilibrium. Since the continuation game in Period t is the same as the original game with a different distribution of the initial asset owners' abilities, so $V_t^{(s)} \in \mathcal{V}$ in every equilibrium. Now in every equilibrium $V_t^{(s)} = \mathcal{T}(V_{t+1}^{(s)})$, so by induction $\mathcal{V} \subset \mathcal{T}^t(\mathcal{V})$ for every positive integer t . The fact that \mathcal{T} is a contraction mapping implies that \mathcal{V} must be a singleton, consisting of the unique fixed point of \mathcal{T} . \square

The key feature of the unique equilibrium is that the asset price is constant over time and is independent of the distribution of the matching quality between current asset owners and their assets. The reason is that for those entrepreneurs who just receive a new asset and those without an asset, the continuation game is the same as the original game, and presence of other entrepreneurs does not change their payoff. Therefore, every period is a fresh start. The above argument makes it clear that this independence is a consequence of the lack of competition among asset owners in either the labor market or the product market. In the labor market, there is an excess supply of homogeneous labor force, and every asset owner can hire a worker at the lowest wage. In the product market, I assume that there is no competition. Nontrivial competition in either market will remove the independence of entrepreneurs' payoff on the distribution of existing matches' qualities, and is likely to remove the independence of price of time and thus to introduce the problem of the optimal time to sell an asset. For that reason, this paper does not study competition in the product market or the labor market.

A consequence of constant asset price is that an entrepreneur who newly acquires an asset either sells it immediately or keeps it forever. When her matching quality with the asset is lower than $a_c^{(s)}$, keeping the asset for one period delivers her a payoff lower than $(1 - \delta)p$, and the price will not change the next period, so she should not wait if she decides to sell her asset. From an asset's perspective, its owner changes every period until it falls into the hand of an entrepreneur with matching quality higher than $a_c^{(s)}$, and then it does not enter the asset market any more. This fact greatly simplifies the characterization of the aggregate output of the economy, which will be taken up in Section 4.

This property of equilibriums with constant asset prices is independent of the contracting environment. Therefore, we will explicitly seek for such an equilibrium in the weak environment, and then try to determine whether every equilibrium has a constant asset price.

2.3.2 The weak contracting environment

When the outcome is not verifiable by a third party, an entrepreneur has to make credible promise on the bonus payment to her employee. Specifically, upon the realization of the good outcome, she must weakly prefer paying the promised bonus to continue her relationship with the employee to renege. By Assumption 1, renege leads to a permanent switch to the static Nash equilibrium between the entrepreneur and all workers, so the entrepreneur will not be able to induce effort from any workers in the future, and has to sell her asset. Hence, the maximum credible bonus payment depends on the resale price of the entrepreneur's asset. This dependence is absent in the strong contracting environment, and makes the analysis in the weak contracting environment more complicated.

In particular, it is not straight forward to establish the uniqueness of equilibrium and the independence of asset price over time in the weak contracting environment. However, we will have an informal discussion of the optimal relational (informal) contract under the assumption of constant asset price, both to generate useful intuition and to set up building blocks in the formal analysis. If the asset price is indeed constant in an equilibrium, the equilibrium will be called *stationary*. Let $V^{(w)}(a, p)$ be the maximum equilibrium payoff of an asset owner with matching quality a when paired with a worker, assuming that the asset price remains constant at p over time. Since neither a nor p changes over time, the value $V^{(w)}(a, p)$ is also independent of time. Assumption 1 implies that every asset owner who did not renege before will receive this payoff.

Assumption 1 also implies that as soon as the entrepreneur reneges, she can induce no worker to exert effort in the future, and thus her best choice is to sell her asset immediately. This implies that the value of the relationship for her is $V^{(w)}(a, p) - p$ to her, and the maximum bonus that she can credibly promise to pay is $\delta[V^{(w)}(a, p) - p]^+$. The value of the relationship for the worker is zero, as he will always get his outside

option². Therefore, $V^{(w)}(a, p)$ solves the following program:

$$\begin{aligned} V^{(w)}(a, p) &= \max_{w, m, e} e(a - m_1) - (1 - e)m_0 - w + \delta V^{(w)}(a, p); \\ \text{s.t. } e &\in \operatorname{argmax}_{\bar{e}} (m_1 - m_0)\bar{e} - c(\bar{e}); \\ em_1 + (1 - e)m_0 + w - c(e) &\geq b; \\ m_z &\leq \delta[V^{(w)}(a, p) - p]^+, \text{ for } z \in \{0, 1\}. \end{aligned}$$

The first and the second constraints are the worker's incentive compatibility constraint and participation constraint, respectively. The third constraint is the bound on the maximum credible bonus payment, which will be called the "no-reneging" constraint. Clearly, the third constraint is never binding for $z = 0$ at optimality, and one can assume WLOG that $m_0 = 0$. Solving for m_1 and w from the first two constraints, one can rewrite the above program as

$$V^{(w)}(a, p) = \max_e ae - c(e) - b + \delta V^{(w)}(a, p); \quad (2.3.5)$$

$$\text{s.t. } c'(e) \leq \delta[V^{(w)}(a, p) - p]^+. \quad (2.3.6)$$

Since $c'(0) \geq 0$, if a positive number $V^{(w)}(a, p)$ solves the above program, it has to be at least as big as p . The maximum number $V^{(w)}(a, p)$ that solves the above program will be taken as the definition of the function $V^{(w)} : [\underline{a}, \bar{a}] \times [0, \infty) \rightarrow [0, \infty)$ evaluated at the point (a, p) , and I put $V^{(w)}(a, p) = 0$ if no positive number $V^{(w)}(a, p)$ satisfies the program Eqs. (2.3.5)-(2.3.6).

Proposition 12. There exists $a_1(p), a_2(p) \in [\underline{a}, \bar{a}]$ such that $a_1(p) \leq a_2(p)$ and the following holds: when $a < a_1(p)$, $V^{(w)}(a, p) = 0$; when $a \in [a_1(p), a_2(p))$, $V^{(w)}(a, p) \in (p, (1 - \delta)^{-1}(g(a) - b))$, and is strictly decreasing in p ; when $a \geq a_2(p)$, $V^{(w)}(a, p) = (1 - \delta)^{-1}(g(a) - b)$.

²I have assumed that incentive payment is made immediately after the realization of z , and the worker is always kept at his outside option as in the beginning of a new relationship. Alternatively, part of the incentive payment can be made through future wages. Levin (2003) shows that this extra option does not increase the entrepreneur's maximum equilibrium payoff.

Proof. It amounts to check that the following values of $a_1(p)$ and $a_2(p)$ have the desired properties, and that $V^{(w)}(a_1(p), p) > p$.

$$\begin{aligned} a_1(p) &= \inf \left\{ a \in [\underline{a}, \bar{a}] : \max_{e \in [0,1]} \frac{1}{1-\delta} [ae - c(e) - b] - \delta^{-1} c'(e) \geq p \right\}; \\ a_2(p) &= \inf \left\{ a \in [\underline{a}, \bar{a}] : a \leq \frac{\delta}{1-\delta} (g(a) - b) - \delta p \right\}. \end{aligned}$$

(Note that by its definition Eq. (2.3.1), g is a convex function, and the inequality in the definition of $a_2(p)$ does not hold for $a = 0$, so that $a \geq a_2(p)$ implies that $a \leq \frac{\delta}{1-\delta} (g(a) - b) - \delta p$.) To see that $V^{(w)}(a_1(p), p) > p$, it suffices to notice that if e is the effort in the optimal relational contract when $a = a_1(p)$, then $a_1(p)e - c(e) - b = (1 - \delta)V^{(w)}(a, p) \geq (1 - \delta)p > 0$, which implies that $V^{(w)}(a, p) - p \geq \delta^{-1}c'(e) > 0$. \square

The proposition implies that when the asset price remains constant at p , asset owners with matching quality lower than $a_1(p)$ will sell their assets immediately, and asset owners with matching quality at least as high as $a_1(p)$ strictly prefers to keep their assets if they did not renege before. Notice that the asset owner's payoff has a discontinuity in a at $a_1(p)$: her payoff will be p if $a < a_1(p)$, but strictly higher than p if $a \geq a_1(p)$. This is because a finite effort has to be induced in every period to sustain cooperation, but the no-renege constraint Eq. (2.3.6) and the assumption that $c'(0) \geq 0$ implies that this requires a strictly positive future surplus $V^{(w)}(a, p) - p$.

Suppose that an asset owner with quality a who did not renege before receives maximum equilibrium payoff $V(a)$ in an relationship, and V is constant over time. Then a new asset owner that did not renege before can decide whether to keep the asset or sell it, and thus receiving payoff $\max\{V(a), p\}$. By Lemma 1, the asset price is $\mathcal{P}(V)$, where $\mathcal{P}(V)$ satisfies

$$\mathcal{P}(V) = \delta \int \max\{V(a), \mathcal{P}(V)\} dF(a). \quad (2.3.7)$$

The above equation uniquely defines map $\mathcal{P} : L^1(F) \rightarrow \mathbb{R}$. Indeed, $\mathcal{P}(V)$ can be obtained by iterating the map $h(p) = \delta \int \max\{V(a), p\} dF(a)$, which is a contraction on \mathbb{R} . Define $\pi : [0, \infty) \rightarrow [0, \infty)$ by

$$\pi(p) = \mathcal{P}(V^{(w)}(\cdot, p)). \quad (2.3.8)$$

Since $V^{(w)}(a, p)$ is weakly decreasing in p for all a and $\mathcal{P}(V)$ is increasing in V , π is a weakly decreasing function. The above discussion implies that if there is a stationary

equilibrium, then the asset price p must be a fixed point of π . Conversely, if p is a fixed point of π , then one can construct an equilibrium in which every asset owner with quality at least $a_1(p)$ uses the optimal policy in Eqs. (2.3.5)-(2.3.6). The following proposition summarizes what we have found so far.

Proposition 13. Assume that F has a probability density. Then π is continuous and has a unique fixed point. In this case, there exists a unique stationary equilibrium in the weak contracting environment that satisfies the following properties: a) the asset price p is the unique fixed point of the function π and is independent of the distribution of the initial asset owners' qualities; b) upon acquiring an asset, an entrepreneur with quality a sells it immediately if $a < a_1(p)$ where a_1 is defined in Proposition 1, and keeps it to receive payoff $V^{(w)}(a, p)$ if $a \geq a_1(p)$.

As expected, the stationary equilibrium in the weak contracting environment also has the desired property that an asset owner either sells her asset immediately or keeps her asset forever in equilibrium. Furthermore, the stationary equilibrium is unique. This uniqueness result is due to the fact that π is weakly decreasing. Intuitively, when the asset price p decreases, every asset owner can promise higher incentive payment and thus receive more profit by keeping her asset, which increases the resale value of an asset. On the other hand, the stationary equilibrium exists when F has a probability density. A counter example can be constructed where F does not have a density and there is no stationary equilibrium.

This proposition characterizes the stationary equilibrium. It does so through a fixed point problem of the mapping $p \mapsto \mathcal{P}(V^{(w)}(\cdot, p))$. However, it is not obvious that every equilibrium will be stationary. This is why we will develop a stronger theorem that says that under some conditions an equilibrium has to be stationary. Before stating that theorem, we remark that the argument that is used to establish the existence and uniqueness of equilibrium in the strong contracting environment does not apply in the weak contracting environment because the counterpart of the operator \mathcal{T} in this environment would not be a contracting in general. Indeed, it is easy to see that the operator \mathcal{T} preserves continuity, and would also be a contraction on $C([\underline{a}, \bar{a}])$ under the uniform norm if it were a contraction on $L^1(F)$, which would imply that its unique fixed point would be a continuous function. However, we have seen in Proposition 1 that in general the value function in the weak contracting environment is not continuous.

Theorem 4. If F has a density and $\pi \circ \pi$ has a unique fixed point, then the stationary equilibrium described in Proposition 2 is the unique equilibrium.

The unique fixed point of π is trivially a fixed point of $\pi \circ \pi$, and the theorem requires that $\pi \circ \pi$ does not have any other fixed points. The proof of the theorem uses a technique similar to the elimination of conditionally dominated strategies found in the literature of bargaining. It constructs a shrinking sequence of intervals that bound the equilibrium asset prices that can arise in equilibrium. The condition that $\pi \circ \pi$ has a unique fixed point implies that the sequence of intervals eventually shrinks to a point.

Before concluding this section, we briefly compare the analyses of the strong and the weak contracting environment. In the strong contracting environment, the value of keeping an asset is independent of the asset price, and therefore the existence and uniqueness of equilibrium can be established under fairly weak conditions. Moreover, the unique equilibrium is stationary. On the other hand, in the weak contracting environment, we add the additional assumption that F has density to establish the existence and uniqueness of the stationary equilibrium, and establishing the uniqueness of equilibrium in general requires an even stronger assumption that $\pi \circ \pi$ has a unique fixed point. In the strong contracting environment, an asset owner's continuation payoff is a continuous function of her matching quality, while in the weak contracting environment there is a discontinuity at $a_1(p)$. The intuition for this discontinuity is that there is a positive feedback in the program Eqs. (2.3.5)-Vw2: a higher future surplus allows the asset owner to promise a higher bonus, which further increases the surplus that can be created from her partnership with the worker. One thing in common between the unique equilibrium in the strong contracting environment and the unique stationary equilibrium in the weak contracting environment is that the asset price in equilibrium is independent of the distribution of the initial asset owners' qualities. Again, this is due to the lack of competition among entrepreneurs in both the labor market and the final product market.

2.4 The Equilibrium Growth Path

The previous section focuses on the partnership game between an asset owner and a worker. This section studies the performance of the economy as a whole over time.

In this section, it will be assumed that F has a density, and in the weak contracting environment we will focus on the unique stationary equilibrium no matter whether the condition in Theorem 2 holds.

The starting point is the observation that in both contracting environments, there is a threshold quality a_c such that upon acquiring an asset, an entrepreneur sells the asset if she discovers that her matching quality with the asset is below a_c , and keeps her asset forever otherwise. In the latter case, an asset owner with quality a produces an output of $y(a)$ every period, where y is strictly increasing in a . In fact, in the strong contracting environment, a_c is the $a_c^{(s)}$ defined in Eq. (2.3.4), and $y(a) = a \min\{1, (c')^{-1}(a)\}$; in the weak contracting environment, a_c is $a_1(p)$ where p is the equilibrium asset price, and $y(a) = a(c')^{-1}(\delta V^{(w)}(a, p) - \delta p)$ if $a \in [a_1(p), a_2(p)]$ and $y(a) = a \min\{1, (c')^{-1}(a)\}$ if $a > a_2(p)$. In what follows, we will also denote the a_c in the weak contracting environment by $a_c^{(w)}$.

Proposition 14. Assume that upon acquiring an asset, an entrepreneur with ability a sells the asset in the next period if $a < a_c$, and keeps the asset to generate output $y(a)$ in every subsequent period if $a \geq a_c$. Then the total output of the economy is

$$Y_t = \int_{a_c}^{\bar{a}} y(a) dF_0(a) + F_0(a_c) \frac{1 - F(a_c)^{t-1}}{1 - F(a_c)} \int_{a_c}^{\bar{a}} y(a) dF(a), \quad (2.4.1)$$

where F_0 is the distribution of the initial asset owners' abilities.

In Eq. (2.4.1), the first term is the contribution from entrepreneurs who initially own assets and decide to keep their assets, and is independent of time. The second term is the contribution from entrepreneurs who acquired assets in the game. This term increases in t , as in each period some entrepreneurs acquire assets, keep them, and produce output in all subsequent periods. Note that $Y_1 = \int_{a_c}^{\bar{a}} y(a) dF_0(a)$, as in the first period only those who initially own assets can produce, and

$$\lim_{t \rightarrow \infty} Y_t = Y_\infty = \int_{a_c}^{\bar{a}} y(a) dF_0(a) + F_0(a_c) \frac{1}{1 - F(a_c)} \int_{a_c}^{\bar{a}} y(a) dF(a). \quad (2.4.2)$$

Therefore, Eq. (2.4.1) can be rewritten as

$$Y_t = Y_1 + [1 - F(a_c)^{t-1}](Y_\infty - Y_1). \quad (2.4.3)$$

In this equation, Y_∞ is the long-run limit of the total output, and $F(a_c)$ characterizes how fast the output converges to its long-run limit. Indeed, Eq. (2.4.3) can be further rewritten as $Y_\infty - Y_t = F(a_c)^{t-1}[Y_\infty - Y_1]$, so the gap between Y_t and Y_∞ shrinks geometrically, and is reduced by a half over approximately $-\ln 2 / \ln(F(a_c))$ periods. Therefore, the smaller $F(a_c)$ is, the faster Y_t converges to Y_∞ . Next I will discuss how Y_1, Y_∞ , and $F(a_c)$ depend on the contracting environment and parameters.

2.4.1 Comparative statics

Proposition 15. Let $p^{(w)}$ be the asset price in the unique stationary equilibrium in the weak contracting environment, and $p^{(s)}$ the asset price in the unique equilibrium in the strong contracting environment. Then $p^{(w)} \leq p^{(s)}$. Moreover, both $p^{(s)}$ and $p^{(w)}$ are strictly decreasing in b and strictly increasing in F in the sense of first order stochastic dominance.

This proposition is not surprising as the asset price is increasing in the asset owners' payoff by keeping their assets, which is decreasing in the unemployment benefit, increasing in their abilities, and lower in the weak contracting environment. A less trivial result is the following.

Proposition 16. In both the strong and the weak contracting environments, the critical level of ability a_c is strictly increasing in the level of unemployment benefit, b , as long as $a_c < \bar{a}$.

There are two effects of b on $a_c^{(s)}$: the direct effect and the indirect effect. The direct effect is that a higher unemployment benefit lowers the surplus that can be created in an employment relationship, thus requires a higher asset owner's ability to enter production. The indirect effect is that a higher unemployment benefit lowers the market price of assets as shown in Proposition 4, and thus reduces the value of the entrepreneur's outside option, which lowers the minimum ability level for production. Proposition 5 says that the direct effect always dominates the indirect effect. Therefore, a higher unemployment benefit would increase the minimum ability level for production.

In the strong contracting environment, the output from each individual entrepreneur $y(a)$ is independent of b . Therefore, Proposition 5 implies that Y_1 is decreasing in b , and $F(a_c)$ is increasing in b , so the convergence to the limit output will be slower. However,

Y_∞ is strictly increasing in a_c and thus in b unless $F_0(a_c) = 0$, as can be directly checked from Eq. (2.4.2). Therefore, though a higher unemployment benefit lowers total output in the economy temporarily, and slows down its convergence to the long-run level, it does increase the long-run level of output by facilitating asset reallocation: more low-ability entrepreneurs are willing to sell their assets now as a result of a higher unemployment benefit.

In the weak contracting environment, for each fixed a , $y(a)$ depends on b . To see this point, define $\hat{V}^{(w)}(a, p) = V^{(w)}(a, p) + (1 - \delta)^{-1}b$. Then Eqs. (2.3.5) and (2.3.6) can be rewritten as

$$\begin{aligned} \hat{V}^{(w)}(a, p) &= \max_e ae - c(e) + \delta \hat{V}^{(w)}(a, p); \\ \text{s.t. } c'(e) &\leq \delta \left[\hat{V}^{(w)}(a, p) - \frac{b}{1 - \delta} - p \right]. \end{aligned}$$

The proof of Proposition 5 implies that $(1 - \delta)^{-1}b + p$ is increasing in b , and thus $\hat{V}^{(w)}(a, p)$ is decreasing in b as the constraint is tighter for a higher b , which implies that the effort level is weakly decreasing in b . Therefore, in the weak contracting environment, $y(a)$ is weakly decreasing in b for all a . Therefore, Y_1 is decreasing in b , but the dependence of Y_∞ on b is ambiguous. Intuitively, in the weak contracting environment, though a higher unemployment benefit facilitates asset reallocation by increasing a_c , it also reduces surplus that can be created in each employment relationship, and thus reduces effort and output in those relationships where the no-renege constraint Eq. (2.3.6) is binding.

It is obvious that for each fixed a , the partnership between an entrepreneur and a worker can generate higher output and higher profit in the strong contracting environment than in the weak contracting environment. However, since the distribution of asset owners' abilities are endogenous and different in the two environments, it is not obvious which environment delivers higher long-run output Y_∞ . Making the output contractible has three effects. First, as shown in Proposition 4, the market price of assets is higher. Secondly, for each fixed a , the output that can be generated by an asset owner with ability a increases. Thirdly, the dependence of a_c on the asset price p changes. Indeed,

I have shown in Eq. (2.3.4) and Proposition 1 that

$$\begin{aligned} g(a_c^{(s)}) &= \max_e a_c^{(s)} e - c(e) = b + (1 - \delta)p^{(s)}; \\ \max_e a_c^{(w)} e - c(e) - \frac{1 - \delta}{\delta} c'(e) &= b + (1 - \delta)p^{(w)}. \end{aligned}$$

Therefore, $a_c^{(s)}$ would be smaller than $a_c^{(w)}$ if $p^{(s)}$ were equal to $p^{(w)}$. Even though $p^{(s)} \geq p^{(w)}$, it is still possible that $a_c^{(s)} < a_c^{(w)}$ in equilibrium, which implies that fewer asset owners are willing to sell their assets in the strong contracting environment, and thus may reduce the long-run output of an economy. The next subsection discusses a "two-type" example that illustrates this possibility.

2.4.2 A "two-type" example

In this subsection, we discuss a "two-type" example that illustrates how the long-run aggregate output depends on the distribution of matching quality a . In this example, the cost of effort $c(e) = \frac{\alpha}{2}e^2$, and the CDF of a is the following:

$$F(a) = \begin{cases} 0, & \text{if } a < a_L - \epsilon; \\ \frac{\mu_L}{2\epsilon}(a - a_L + \epsilon), & \text{if } a \in [a_L - \epsilon, a_L + \epsilon]; \\ \mu_L, & \text{if } a \in [a_L + \epsilon, a_H]; \\ 1, & \text{if } a \geq a_H. \end{cases}$$

Here $\mu_L \in (0, 1)$, $0 < \epsilon < a_L < a_H - \epsilon$. In other words, F consists of a point mass of $(1 - \mu_L)$ on a_H and a mass μ_L uniformly distributed on $[a_L - \epsilon, a_L + \epsilon]$. Call asset owners with quality a_H the high-quality owners, and other asset owners low-quality owners. We are interested in the limit $\epsilon \rightarrow 0$, and the reason to introduce this continuum is to ensure the existence of the stationary equilibrium in the weak contracting environment. Let $\mu_H = 1 - \mu_L$. The following assumption will be maintained:

Assumption 8. $\frac{1}{\delta}\alpha + \frac{1}{2}\alpha + b < a_H < \frac{1 - \delta\mu_L}{2\delta\mu_H}(\alpha - 2b) + \frac{1}{2}\alpha + b$.

The goal is to determine how the stationary equilibrium and the long-run aggregate output depends on a_L in the limit $\epsilon \rightarrow 0$. Clearly, all high quality asset owners keep their assets in any stationary equilibrium. It remains to find whether low-quality asset owners keep their assets. Let $q_\epsilon(a_L)$ be the fraction of low-quality asset owners who sell

their assets, and

$$q_{0+}(a_L) = \lim_{\epsilon \rightarrow 0} q_\epsilon(a_L).$$

The following proposition characterizes the limit of $q_\epsilon(a_L)$ for $\epsilon \rightarrow 0$ in the strong contracting environment:

Proposition 17. Consider the strong contracting environment under Assumption 2. Then an employee working for a high-quality asset owner exerts effort one. Let

$$a_{LS} = \sqrt{\frac{\delta\mu_H(2a_H - \alpha)\alpha + 2(1 - \delta)\alpha b}{1 - \delta\mu_L}}. \quad (2.4.4)$$

Then $q_{0+}(a_L) = 1$ if $a_L < a_{LS}$, and $q_{0+}(a_L) = 0$ if $a_L > a_{LS}$.

The proposition says that $q_{0+}(a_L)$ is not continuous in a_L . In fact, in this case $q_{0+}(a_L)$ is equal to the indicator that the low-quality asset owners sell their assets when $\epsilon = 0$. In that case, the asset price is $p = \delta\mu_H V_H / (1 - \delta\mu_L)$ when all low-quality asset owners sell their assets, and it is strictly optimal for a low-quality asset owner to do so when $p > \max_e (a_L e - \frac{\alpha}{2} e^2 - b)$. Assumption 2 implies that the optimal choice of e is less than one for a_L in a neighborhood of a_{LS} , and the inequality holds if and only if $a_L < a_{LS}$. In the strong contracting environment, the asset price changes continuously with q , and a low-quality asset owner is indifferent between keeping and selling her asset when $a_L = a_{LS}$, so she prefers to keep her asset when $a_L > a_{LS}$.

In the weak contracting environment, we have the following proposition:

Proposition 18. Consider the weak contracting environment under Assumption 2. Then an employee working for a high-quality asset owner exerts effort one. Let

$$a_{LW1} = \frac{1 - \delta}{\delta} \alpha + a_{LS}; \quad (2.4.5)$$

$$a_{LW2} = \frac{1 - \delta}{\delta(1 - \delta\mu_L)} \alpha + a_{LS}. \quad (2.4.6)$$

Then $q_{0+}(a_L) = 1$ if $a_L \leq a_{LW1}$, $q_{0+}(a_L) = 0$ if $a_L \geq a_{LW2}$, and q_{0+} is linear on $[a_{LW1}, a_{LW2}]$.

In this case, q_{0+} is a continuous function. In particular, q_{0+} is strictly between zero and one for a_L between a_{LW1} and a_{LW2} : for these values of a_L only a fraction of low-quality asset owners sell their assets. Furthermore, an asset owner's payoff is discontinuous in

her quality in this case: her payoff jumps up by a finite amount when her quality exceeds the threshold $a_1(p)$ in Proposition 1, and she strictly prefers keeping her asset when her quality equals the threshold. This discontinuity is related to the fact that there does not exist a stationary equilibrium when $\epsilon = 0$ and $a_L \in (a_{LW1}, a_{LW2})$. The proof of the proposition is the application of the general Proposition 1 to this special case.

As expected, both a_{LW1} and a_{LW2} are bigger than a_{LS} , which means that low-quality asset owners are more likely to sell their assets in the weak contracting environment. However, the difference between a_{LW2} and a_{LS} vanishes as δ approaches one. This is also intuitive: the lack of formal contracts is not important when all agents are sufficiently patient.

Figure 1 shows how $q_{0+}(a_L)$ depends on a_L in both contracting environments in a numerical simulation. In this simulation, $a_H = 2$, $\delta = 0.8$, $\alpha = 1$, $b = 0$, and $\mu_L = 0.9$. It is straight forward to compute that $a_{LS} = 0.926$, $a_{LW1} = 1.176$, and $a_{LW2} = 1.819$. As expected, q_{0+} is not continuous in the strong contracting environment, and is linear on $[a_{LW1}, a_{LW2}]$ in the weak contracting environment.

Figure 2 shows how the long-run aggregate output, Y_∞ , depends on a_L , where it is assumed that $F_0 = F$, so the initial distribution of matching quality is the same as the population distribution. The first thing to notice is that Y_∞ is not monotonic in a_L : when a_L is small, low-quality asset owners sell their assets, which leads to a high long-run output; when low-quality asset owners keep their assets, the long-run output is increasing in their quality.

It is interesting to compare the long-run aggregate outputs in the two contracting environment. When $a_L \leq a_{LS}$, all low-quality asset owners sell their assets in both contracting environment, and Assumption 2 implies that the no-reneging constraint in the weak contracting environment is not binding for high-quality asset owners, so long-run outputs are the same in both environments. When $a_L \in (a_{LS}, a_{LW1}]$, all low-quality asset owners sell their assets in the weak contracting environment, and all of them keep their assets in the strong contracting environment, so long-run output is higher in the weak contracting environment. When $a_L \in (a_{LW1}, a_{LW2}]$, some low-quality asset owners keep their assets in the weak contracting environment, and the no-reneging constraint is binding for them, so their outputs are lower than in the first best. An increase in a_L has two effects on Y_∞ in this region. First, more low-quality asset owners keep their




FIGURE 2.1: Fraction of low-quality asset owners who sell their assets. This is in the limit $\epsilon \rightarrow 0$. Values of parameters are $a_H = 2$, $\delta = 0.8$, $\alpha = 1$, $b = 0$, and $\mu_L = 0.9$.




FIGURE 2.2: Long-run aggregate output as a function of a_L .

assets in the weak contracting environment. This will be called the "composition effect". Secondly, those asset owners who keep their assets generate higher outputs. This will be called the "level effect". The composition effect lowers the long-run aggregate output, while the level effect raises it. In the strong contracting environment, Y_{∞} is increasing in a_L as long the level effect is present. In the weak contracting environment, Y_{∞} is not monotonic in a_L as the two effects act simultaneously. Since the low-quality asset

owners who keep their assets are less efficient in the weak environment, the aggregate output is lower in the weak environment when $q_{0+}(a_L)$ is close to zero. Finally, when $a_L > a_{LW2}$, all low-quality asset owners keep their assets in both environments. The long-run aggregate output in the weak contracting environment is weakly lower than in the strong contracting environment, and strictly so when the no-renege constraint is binding for the low-quality asset owners. Therefore, if the initial distribution of assets is not efficient in a developing country, and a_L is between a_{LS} and a_{LW1} , the country will have a higher long-run GDP in the weak contracting environment.

Hsieh and Klenow (2009) documents the improvement of asset distribution in China from 1998 to 2005. It turns out that the labor legislation was weak in China during that period. The new Labour Contract Law took effect in 2008, and the old Labour Law was in effect from 1995 to 2007. The old Labour Law had a few weak points that had been exploited by employers to extract rents from employees. For example, every contract can contain a probation period, and an employer can sign a short-term contract with an employee with a relatively long probation period and fire the employee after the probation period. The new Labour Contract Law fixes this issue and others, and strengthens the contracting environment. However, the improvement of asset allocative efficiency documented in Hsieh and Klenow (2009) happened during the era of the old Labour Law.

2.5 A Continuum of Contracting Environment

In the discussion so far, I have considered two extreme contracting environment: the strong contracting environment in which the first best can be achieved in every matching of an asset owner and a worker, and the weak contracting environment in which no enforceable incentive contract can be written. This section extends the model to allow for a continuum of contracting environment indexed by a real number θ , which measures the strength of the contracting environment. This extension both builds a bridge between the two extreme contracting environment and helps to flesh out the economic meaning of the measurability of output.

This extension is based on the multi-tasking moral hazard model proposed by Baker (1992) and Baker (2002). In each period, an imperfect performance measure is generated

and can be verified by the court, where "imperfect" means that it is not always equal to the revenue generated for the entrepreneur. Correspondingly, the worker can choose an action ϵ that does not generate revenue, but boosts the performance measure, in addition to the normal effort e . Specifically, I replace steps 6 and 7 in the timing described in Section 2 by the following:

- 6'**. If a wage offer is accepted, the wage w is paid, and the worker chooses effort $e \in [0, \bar{e}]$ and manipulation action $\epsilon \in [0, \bar{\epsilon}]$ at a private cost $c(e) + c_\epsilon(\epsilon)$. I assume that both c and c_ϵ are strictly increasing, twice continuously differentiable and have strictly positive second derivatives. I also assume that $c(0) = c_\epsilon(0) = c'(0) = c'_\epsilon(0) = 0$.
- 7'**. Within each pair of matched entrepreneur and worker, an output $z = (z_r, z_f) \in \{0, 1\}^2$ is realized and observed publicly, where z_r and z_f are independent conditional on (e, ϵ) , $\Pr(z_r = 1|e, \epsilon) = e$, and $\Pr(z_f = 1|e, \epsilon) = e \cos \theta + \epsilon \sin \theta$. The entrepreneur receives cash az_r , and a bonus payment is made. To make sure that all probabilities are between 0 and 1 for all θ , I assume that $\bar{e}^2 + \bar{\epsilon}^2 \leq 1$.

I assume that z_r is not verifiable by a third party, but z_f is. Therefore, the incentive scheme $m(z_r, z_f)$ has the following interpretation: the incentive contract specifies that the entrepreneur pays the worker bonus $m(0, 0)$ when $z_f = 0$ and $m(0, 1)$ when $z_f = 1$; in addition, the entrepreneur promises to pay the worker $m(1, z_f) - m(0, z_f)$ if $z_r = 1$, which she can choose whether to deliver. Though z_f does not measure the entrepreneur's revenue, it can be useful in providing incentives, because the maximum bonus payment that the entrepreneur can credibly promise to pay associated with z_r is bound by the future surplus, while there is no such bound for bonus associated with z_f .

Here $\theta \in [0, \pi/2]$ is the characteristic of the contracting environment. When θ is small, the correlation between z_r and z_f is high, and z_f can become a useful incentive instrument. When θ is big, z_f mainly depends on the manipulation action ϵ , which is socially costly. Therefore, the smaller θ is, the stronger the contracting environment is. In particular, $\theta = 0$ and $\theta = \pi/2$ are equivalent to the strong contracting environment and the weak contracting environment discussed in Section 2, respectively. In what follows I will assume that $\theta \in (0, \pi/2)$.

To analyze the equilibrium in this multi-tasking environment, first consider the asset owner's payoff when she uses purely formal contracts, so that $m(z_r, z_f)$ is independent

of z_r . Then renegeing is not possible, and the asset owner's payoff from production $V_f(a, \theta)$ is independent of the asset price:

$$(1 - \delta)V_f(a, \theta) = \max_{e, \epsilon, m} ae - c(e) - c_\epsilon(\epsilon) - b;$$

$$s.t. (e, \epsilon) \in \operatorname{argmax}_{(\bar{e}, \bar{\epsilon})} m(\bar{e} \cos \theta + \bar{\epsilon} \sin \theta) - c(\bar{e}) - c_\epsilon(\bar{\epsilon}).$$

The constraint in the above program is the worker's incentive compatibility constraint, and I have used his participation constraint to eliminate the wage w and rewrite the entrepreneur's objective function. When $\theta > 0$, the first best cannot be achieved: though the entrepreneur is free to choose incentives of arbitrary strength, any incentive would induce a positive manipulation ϵ , which is a social waste.

Let $V(a, p, \theta)$ be the maximum equilibrium continuation payoff that an asset owner with ability a can receive by keeping her asset, when the asset price remains constant at p over time. Then $V(a, p, \theta)$ is the maximum solution to the following program:

$$V(a, p, \theta) = \max_{e, \epsilon, m} ae - c(e) - c_\epsilon(\epsilon) - b + \delta V(a, p, \theta);$$

$$s.t. (e, \epsilon) \in \operatorname{argmax}_{(\bar{e}, \bar{\epsilon})} (e \cos \theta + \epsilon \sin \theta)[e(m(1, 1) - m(1, 0)) + (1 - e)(m(0, 1) - m(0, 0))] +$$

$$+ em(1, 0) + (1 - e)m(0, 0) - c(\bar{e}) - c_\epsilon(\bar{\epsilon});$$

$$m(1, z_f) - m(0, z_f) \leq \delta[V(a, p, \theta) - \max\{V_f(a, \theta), p\}]^+, \text{ for } z_f \in \{0, 1\}.$$

The first constraint is the worker's incentive compatibility constraint, and the second constraint is the entrepreneur's no-renegeing constraint. There are two differences between this constraint and Eq. (2.3.6). First, this constraint has to be satisfied for both values of z_f , which means that the bonus promise made by the entrepreneur has to be credible for both values of z_f . Secondly, on the right hand side p in Eq. (2.3.6) is replaced by $\max\{V_f(a, \theta), p\}$, since after renegeing the entrepreneur has the option to keep her asset and use purely formal contracts. I will still follow the convention that $V(a, p, \theta) = 0$ if the above program does not have a positive solution.

Lemma 9. Suppose that $V(a, p, \theta) > 0$ for some (a, p, θ) . Then there is always an optimal incentive scheme m for (a, p, θ) such that $m(0, 0) = 0$, and $m(1, 1) - m(0, 1) = m(1, 0) - m(0, 0)$.

Using this lemma, the program of $V(a, p, \theta)$ can be rewritten in terms of $m_r = m(1, 1) - m(0, 1) = m(1, 0) - m(0, 0)$ and $m_f = m(0, 1) - m(0, 0)$:

$$V(a, p, \theta) = \max_{e, \epsilon, m_r, m_f} ae - c(e) - c_\epsilon(\epsilon) - b + \delta V(a, p, \theta); \quad (2.5.1)$$

$$s.t. \quad (e, \epsilon) \in \operatorname{argmax}_{(\bar{e}, \bar{\epsilon})} m_f(\bar{e} \cos \theta + \bar{\epsilon} \sin \theta) + m_r e - c(\bar{e}) - c_\epsilon(\bar{\epsilon}) \quad (2.5.2)$$

$$m_r \leq \delta[V(a, p, \theta) - \max\{p, V_f(a, \theta)\}]^+. \quad (2.5.3)$$

Proposition 19. Let $\pi_\theta(p) = \mathcal{P}(V(\cdot, p, \theta))$ where \mathcal{P} is defined in Eq. (2.3.7). Then Proposition 2 and Theorem 2 hold with π replaced by π_θ .

This proposition implies that the analysis of the equilibrium growth path in Section 4 applies to every contracting environment θ . In particular, for every θ there exists $a_c(\theta) \in [\underline{a}, \bar{a}]$ such that asset owners with abilities $a < a_c(\theta)$ sell their assets immediately, and asset owners with $a > a_c(\theta)$ keep their assets for ever. The final proposition in this section shows that purely formal contracts almost never exists in equilibrium.

Proposition 20. Fix $\theta > 0$. Then in the unique stationary equilibrium, for all but at most one $a \in (a_c(\theta), \bar{a}]$, every entrepreneur with ability a relies on informal incentives (i.e. chooses $m_r > 0$).

Intuitively, since informal incentive ties the agent's compensation directly to the principal's revenue, it is a superior incentive device than the formal contract. Therefore, in generic case the principal always adds an informal component to the compensation scheme. This proposition gives a reason why informal incentives prevail in the real world. This does not contradict the intuition that the outcome should converge to "fully formal contract" when $\theta \rightarrow 0$. As $\theta \rightarrow 0$, $V_f(a, \theta)$ converges to the first-best payoff, and feasibility means that $V(a, p, \theta) - \max\{p, V_f(a, \theta)\}$ is zero or close to zero, and Eq. (2.5.3) implies that the informal component of the incentive, m_r , is small.

2.6 Conclusion

In this paper, we study the market of assets when the value of each asset is generated through an employment relationship. The reallocation of assets can become a temporary source of economic growth when the initial allocation of assets is not efficient. By

comparing two contracting environments and the unique stationary equilibrium for different values of the unemployment benefit, we show that some factors that negatively affect output and surplus in a single relationship may benefit the whole economy when reallocation of assets is ongoing. This paper leaves out the interaction between asset reallocation and competition among businesses in the labor and final product market, which can be interesting and is left for future research.

2.7 Appendix

Proof of Lemma 1. Since the sellers own all the bargaining power, each seller will offer her buyer's continuation value when the buyer refuses to buy in the current period, but this continuation value is independent of the seller's own strategy. This proves the first claim. Let p_t be the market price of assets at period t and V_t the value of an asset for its buyer, which does not depend on the identity of the asset due to Assumption 1 and does not depend on the identity of the buyer either since without owning assets all entrepreneurs are the same. Clearly, $V_t \leq \max_{e \in [0,1]} (1 - \delta)^{-1} (\bar{e} - c(e)) < \infty$, the entrepreneur's first-best payoff when she has the maximum ability \bar{a} . Let $\Delta = \sup_t (V_t - p_t)$. Then in each transaction, the buyer's continuation value without buying the asset is at most $\delta^T \Delta$, if she successfully buys an asset T periods later, which is assumed to happen with probability q_T when she uses the optimal strategy. The buyer's continuation value is at most $\sum_{T=1}^{\infty} q_T \delta^T \Delta \leq \delta \Delta$. Since the seller has all the bargaining power, she offers the price of at least $p_t = V_t - \delta \Delta$. This holds for all t , so $\Delta = \sup_t (V_t - p_t) \leq \delta \Delta$, which implies that $\Delta = 0$. The proof that workers receive no surplus is similar.

Proof of Proposition 2. We have seen that a fixed point of π gives rise to a stationary equilibrium, and in a stationary equilibrium the asset price has to be a fixed point of π . It remains to show that π is continuous and has a unique fixed point. According to Proposition 1 and Eq. (2.3.7),

$$\pi(p) = \delta \int_{a_1(p)}^{\bar{a}} V^{(w)}(a, p) dF(a) + \delta \pi(p) F(a_1(p)).$$

Therefore,

$$\pi(p) = \frac{\delta}{1 - \delta F(a_1(p))} \int_{a_1(p)}^{\bar{a}} V^{(w)}(a, p) dF(a). \quad (2.7.1)$$

Let

$$\gamma(a) = \max_{e \in [0,1]} (1 - \delta)^{-1} [ae - c(e)] - \delta^{-1} c'(e). \quad (2.7.2)$$

Then γ is a weakly increasing and convex function, since its objective function is strictly increasing and linear in a for $e > 0$. Let $a_0 = \sup\{a : \gamma(a) \leq 0\}$. Then for $a > a_0$, the optimal effort is positive, and thus γ is strictly increasing. In what follows I restrict the domain of γ to $[a_0, \infty)$. The proof of Proposition 1 implies that $a_1(p) = \gamma^{-1}(p + (1 - \delta)^{-1}b)$, so a_1 is a strictly increasing and concave function. In particular, a_1 is

a continuous function. Define $\xi(a, x) = (1 - \delta)^{-1}[a(c')^{-1}(\delta x) - c((c')^{-1}(\delta x))]$. Then for $a \in [a_1(p), a_2(p)]$, $V^{(w)}(a, p) = \xi(a, V^{(w)}(a, p) - p)$. Since c is twice continuously differentiable and $c''(x) > 0$ for all $x \in [0, 1]$, $\xi(a, x)$ is continuously differentiable in x . It is also easy to see that $\partial\xi/\partial x$ is strictly increasing in a . The implicit function theorem implies that unless $\partial\xi/\partial x(a, V^{(w)}(a, p) - p) = 1$, $V^{(w)}(a, p)$ is continuously differentiable in p . However, $\partial\xi/\partial x(a, V^{(w)}(a, p) - p) = 1$ when $a = a_1(p)$. Therefore, $V^{(w)}(a, p)$ is continuously differentiable in p for $a \in (a_1(p), a_2(p))$. For $a > a_2(p)$, $V^{(w)}(a, p)$ is independent of p . The continuity of π follows from Eq. (2.7.1), the continuity of F , a_1 , and $V(a, \cdot)$, and the dominated convergence theorem.

Now $\pi(0) > 0$ and $\pi(p) = \delta p < p$ for $p > (1 - \delta)^{-1}(g(\bar{a}) - b)$. Therefore, π has a fixed point in $(0, (1 - \delta)^{-1}(g(\bar{a}) - b))$. Since π is weakly decreasing, the fixed point is unique.

Proof of Theorem 2. Let p_t be the asset price in Period t . Then after acquiring an asset at period t , an entrepreneur can always sell the asset in the next period, receiving proceeds δp_{t+1} . Therefore, the value of an asset in Period t is at least δp_{t+1} , and by Lemma 1, $p_t \geq \delta p_{t+1}$ in every equilibrium.

It is clear that an asset owner's optimal decision on whether to sell her asset does not depend on the realized outcomes in the past, unless she reneged before. Suppose that in an equilibrium it is optimal for an asset owner with ability a to sell her asset t periods after she acquires the asset. Then the future surplus is zero for her $(t - 1)$ periods after her acquisition of the asset, and thus she cannot credibly promise any bonus payment in the $(t - 1)$ th period since acquisition. Consequently, there cannot be any output in the $(t - 1)$ th period since her acquisition, and it must be optimal for her to sell her asset in the $(t - 1)$ th period. (Waiting for one more period cannot strictly improve her payoff as I have shown that $p_{t-1} \geq \delta p_t$ in equilibrium. By induction, it must be optimal for her to sell her asset immediately. Therefore, upon acquiring an asset, an entrepreneur either keeps the asset for ever or sells the asset immediately after observing her ability. I refer to her payoff in the former case by her payoff from by keeping her asset. I first establish two lemmas.

Lemma 10. If the asset price is at most (at least) p in every period in an equilibrium, then the continuation payoff of an asset owner with ability a by keeping her asset is at least (at most, respectively) $V^{(w)}(a, p)$ starting from every period

Proof. Assume that the asset price is at most p in every period in an equilibrium. Then every entrepreneur with ability $a \geq a_1(p)$ (which is defined in Proposition 1) has the option to adopt the optimal policy in the solution to the program Eqs. (2.3.5)-(2.3.6) that defines $V^{(w)}(a, p)$ and thereby receive payoff $V^{(w)}(a, p)$, and the no-renege constraint is satisfied in every period since a lower p only relaxes that constraint. Therefore, an asset owner with ability a will receive continuation payoff of at least $V^{(w)}(a, p)$ by keeping her asset, as long as she did not renege before.

Now assume instead that the asset price is at least p in every period in an equilibrium. Suppose that an asset owner with ability a earns continuation payoff $v > V^{(w)}(a, p)$ by keeping her asset starting from some period. Then adopting her strategy when the asset price is constantly at p would satisfy all the IC, IR and no-renege constraints as a lower asset price relaxes the no-renege constraints, so $V^{(w)}(a, p) \geq v$ by the construction of $V^{(w)}(a, p)$, a contradiction. Therefore, an asset owner with ability a can earn at most continuation payoff $V^{(w)}(a, p)$ by keeping her asset starting from every period. \square

Lemma 11. If in an equilibrium, the continuation payoff of an asset owner with ability a by keeping her asset is bounded from above (below) by $V(a)$ for F -almost every $a \in [\underline{a}, \bar{a}]$ starting from every period, and $V \in L^1(F)$, then the asset price is at most (at least, respectively) $\mathcal{P}(V)$.

Proof. Let $V_t(a)$ be the continuation payoff of an asset owner with ability a by keeping her asset starting from period t . Then $V_t(a) \in [0, V^{(w)}(a, 0)]$ for all a and t by the previous lemma. Let p_t be the asset price in period t . Then upon acquiring an asset at period t , an entrepreneur with ability a can either keep her asset and receives continuation payoff $\delta V_{t+1}(a)$ or sell her asset in the next period and receives payoff δp_{t+1} . Therefore, $p_t = h(V_{t+1}, p_{t+1})$ for all t , where

$$h(\tilde{V}, p) = \delta \int \max\{\tilde{V}(a), p\} dF(a), \text{ for } \tilde{V} \in L^1(F), p \in [0, \infty).$$

By construction $\mathcal{P}(V)$ is a fixed point of $h(V, \cdot)$. It is easy to see that h is weakly increasing in p , and h is also weakly increasing in \tilde{V} in the sense that if $\tilde{V}'(a) \geq \tilde{V}''(a)$ for F -almost every $a \in [\underline{a}, \bar{a}]$, then $h(\tilde{V}', p) \geq h(\tilde{V}'', p)$ for all $p \in [0, \infty)$. Moreover, for every $\tilde{V} \in L^1(F)$, $|h(\tilde{V}, p') - h(\tilde{V}, p'')| \leq \delta |p' - p''|$ for all $p', p'' \in [0, \infty)$.

Now suppose that $V_t(a) \geq V(a)$ for F -almost every $a \in [\underline{a}, \bar{a}]$ and all t , but $\Delta = \sup_t(\mathcal{P}(V) - p_t) > 0$. Then for every t ,

$$\begin{aligned}
& \mathcal{P}(V) - p_t \\
&= h(V, \mathcal{P}(V)) - h(V_{t+1}, p_{t+1}) \\
&= [h(V, \mathcal{P}(V)) - h(V, p_{t+1})] + [h(V, p_{t+1}) - h(V_{t+1}, p_{t+1})] \\
&\leq \delta(\mathcal{P}(V) - p_{t+1})^+ + 0 \\
&\leq \delta\Delta.
\end{aligned}$$

This holds for all t , so $\Delta = \sup \mathcal{P}(V) - p_t \leq \delta\Delta$. This implies that $\Delta \leq 0$ or $\Delta = \infty$. However, $\mathcal{P}(V) \in [0, \delta V(\bar{a})]$ and $p_t \geq [0, \delta(1 - \delta)^{-1}(g(\bar{a}) - b)]$ for all t , so $\Delta \leq 0$, a contradiction. The other half of the lemma can be proved in the same manner. \square

Now I return to the proof of Theorem 2. Let $l_0 = 0$, $u_0 = \mathcal{P}(V^{(w)}(\cdot, l_0)) = \pi(l_0)$ where π is defined in Eq. (2.3.8), and recursively define $l_n = \pi(u_{n-1})$ and $u_n = \pi(l_n)$ for all positive integers n . Then using the above two lemmas and by induction, one can show that $p_t \in [l_n, u_n]$ for all n in every equilibrium. The monotonicity of π and the fact that $l_0 = 0 \leq l_1$ implies that (l_n) is a weakly increasing sequence, while (u_n) is a weakly decreasing sequence. Let $l_\infty = \lim_{n \rightarrow \infty} l_n$, and $u_\infty = \lim_{n \rightarrow \infty} u_n$. Then the continuity of π implies that $l_\infty = \mathcal{P}(V^{(w)}(\cdot, u_\infty))$, and $u_\infty = \mathcal{P}(V^{(w)}(\cdot, l_\infty))$. Therefore, both l_∞ and u_∞ are fixed points of $\pi \circ \pi$. Consequently, $l_\infty = u_\infty$, and an equilibrium has to be stationary.

Proof of Proposition 3. Let m_t be the total mass of assets sold in Period t . In every period t , the entrepreneurs who newly acquire assets are randomly drawn from the distribution F , so the distribution of these entrepreneurs' abilities is given by F . Among them, those with abilities lower than a_c sell their asset in the next period, and each entrepreneur either sell her asset immediately or keeps her asset forever, so $m_{t+1} = F(a_c)m_t$. The same argument implies that $m_1 = F_0(a_c)$. Therefore, by induction $m_t = F_0(a_c)F(a_c)^{t-1}$. The distribution of the abilities of those entrepreneurs who acquire assets in Period $t > 1$ and decide to keep their assets is F conditional on the event that $a \geq a_c$, so the distribution function is $\frac{1}{1-F(a_c)}F(a)1_{a \geq a_c}$. By induction, this is also the distribution function of the abilities of all asset owners in period t who did not own assets at the

beginning, for all t . Their total mass is $m_1 - m_t = F_0(a_c)[1 - F(a_c)^{t-1}]$, and thus their total output is $F_0(a_c)[1 - F(a_c)^{t-1}][1 - F(a_c)]^{-1} \int_{a_c}^{\bar{a}} y(a) dF(a)$. The initial asset owners who have abilities at least as high as a_c keep their assets and produce total output $\int_{a_c}^{\bar{a}} y(a) dF_0(a)$ in every period. This proves Eq. (2.4.1).

Proof of Proposition 4. In the strong contracting environment, each entrepreneur with ability $a \geq a_1(p^{(w)})$ can choose to keep her asset for ever and receive payoff $(1 - \delta)^{-1}(g(a) - b)$, which is at least as big as $V^{(w)}(a, p^{(w)})$ because it is the value of the program Eqs. (2.3.5)-Vw2 without the constraint. Therefore, $V^{(s)}(a) \geq V^{(w)}(a, p^{(w)})$ for $a \geq a_1(p^{(w)})$. The inequality also holds for $a < a_1(p^{(w)})$ as $V^{(w)}(a, p) = 0$ in that case. It is easy to see that $p^{(s)} = \mathcal{P}(V^{(s)})$. That $\mathcal{P}(V)$ is the limit of the iteration of the contraction mapping $p \mapsto \delta \int \max\{V(a, p)\} dF(a)$ implies that $p^{(s)} = \mathcal{P}(V^{(s)}) \geq \mathcal{P}(V^{(w)}(\cdot, p^{(w)})) = p^{(w)}$.

Let $\hat{V}(a) = (1 - \delta)[g(a) - b]$. Then $V^{(s)}(a) = \hat{V}(a)$ for $a \geq a_c^{(s)}$. Therefore, $\delta \int \max\{\hat{V}(a), p^{(s)}\} dF(a) = \delta \int \max\{V^{(s)}(a), p^{(s)}\} dF(a) = p^{(s)}$, which means that $\mathcal{P}(\hat{V}) = p^{(s)}$. Since \hat{V} is strictly decreasing in b , $p^{(s)}$ is also strictly decreasing in b . Finally, $p^{(w)}$ is strictly decreasing in b because $V^{(w)}(a, p)$ is strictly decreasing in b , and $p^{(w)}$ is the unique fixed point of the map $p \mapsto \mathcal{P}(V^{(w)}(\cdot, p))$. Monotonicity with respect to F can be proved in the same way.

Proof of Proposition 5. First consider the strong contracting environment. By the proof of the previous proposition, $p^{(s)} = \mathcal{P}(\hat{V})$ where $\hat{V}(a) = (1 - \delta)^{-1}[g(a) - b]$. Therefore, $p^{(s)}(b) = \frac{\delta}{1 - \delta} \int_{a_c^{(s)}(b)}^{\bar{a}} [g(a) - b] dF(a) + \delta F(a_c^{(s)}(b)) p^{(s)}(b)$, where I have made the dependence of b explicit. Differentiating both sides yields

$$\frac{dp^{(s)}}{db} = -\frac{\delta(1 - F(a_c^{(s)}(b)))}{1 - \delta} + \delta F(a_c^{(s)}(b)) \frac{dp^{(s)}}{db} - \delta F'(a_c^{(s)}(b)) \frac{da_c^{(s)}}{db} [(1 - \delta)^{-1}(g(a_c^{(s)}(b)) - b) - p^{(s)}(b)],$$

but the last term is zero by Eq. (2.3.4). Therefore,

$$\frac{dp^{(s)}}{db} = -\frac{\delta[1 - F(a_c^{(s)}(b))]}{(1 - \delta)[1 - \delta F(a_c^{(s)}(b))]} > -\frac{1}{1 - \delta}.$$

Therefore, $(1 - \delta)p^{(s)}(b) + b$ is strictly increasing in b , which implies that $a_c^{(s)}$ is strictly increasing in b since $g(a_c^{(s)}) = (1 - \delta)p^{(s)} + b$.

Now consider the weak contracting environment. Let $\alpha = dp^{(w)}/db$. Suppose that $\alpha \leq -1/(1-\delta)$. By the proof of Proposition 1, $\gamma(a_1) = p + (1-\delta)^{-1}b$ where γ is defined in Eq. (2.7.2). The inverse function theorem implies that $\gamma'(a_1) \frac{da_1}{db} = \alpha + (1-\delta)^{-1}$. It is easy to see from Eq. (2.7.2) that γ is a convex function, and the envelope theorem implies that $\gamma'(a_1) = (1-\delta)^{-1}e > 0$. Therefore,

$$\frac{da_1}{db} = \frac{\alpha + (1-\delta)^{-1}}{\gamma'(a_1)}. \quad (2.7.3)$$

Defining $\hat{V}^{(w)}(a, p) = V^{(w)}(a, p) + (1-\delta)^{-1}b$, the program Eqs. (2.3.5)-(2.3.6) can be rewritten as

$$\begin{aligned} \hat{V}^{(w)}(a, p) &= \max_e ae - c(e) + \delta \hat{V}^{(w)}(a, p); \\ \text{s.t. } c'(e) &\leq \delta \left[\hat{V}^{(w)}(a, p) - \frac{b}{1-\delta} - p \right]. \end{aligned}$$

Consider the case where $p = p^{(w)}(b)$. When $\alpha \leq -1/(1-\delta)$, the constraint becomes looser for a higher b , and therefore $\hat{V}^{(w)}(a, p^{(w)}(b))$ is increasing in b , which implies that $dV^{(w)}(a, p^{(w)}(b))/db \geq -(1-\delta)^{-1}$ for all $a > a_1(p)$. Now differentiating both sides of the equation

$$p^{(w)}(b) = \delta \int_{a_1}^{\bar{a}} V^{(w)}(a, p^{(w)}(b)) dF(a) + \delta F(a_1) p^{(w)}(b)$$

yields that

$$\alpha \geq \frac{\delta[1 - F(a_1)]}{1-\delta} + \delta F(a_1)\alpha - \delta f(a_1) \frac{da_1}{db} [V^{(w)}(a_1, p^{(w)}) - p^{(w)}],$$

where f is the probability density of the distribution F . Substituting in Eq. (2.7.3) yields that

$$\alpha \geq \frac{\delta[1 - F(a_1)]}{1-\delta} + \delta F(a_1)\alpha - \frac{1}{\gamma'(a_1)} \delta f(a_1) [1 + (1-\delta)\alpha] [V^{(w)}(a_1, p^{(w)}) - p^{(w)}].$$

The last term is non-negative since $\alpha \leq -1/(1-\delta)$, and $\alpha < \frac{\delta[1-F(a_1)]}{1-\delta} + \delta F(a_1)\alpha$ when $\alpha \leq -1/(1-\delta)$, a contradiction. Therefore, $\alpha > -1/(1-\delta)$, and thus a_1 is increasing in b by Eq. (2.7.3).

Proof of Proposition 6. For an employee working for a high-quality asset owner, the optimal effort is $\min\{a_H/\alpha, 1\}$, which equals one by Assumption 2. Let

$$V_H = \frac{1}{1-\delta} \left(a_H - \frac{\alpha}{2} - b \right), \quad (2.7.4)$$

the payoff of a high-quality asset owner. Suppose that all low-quality asset owners sell their assets. Then the asset price p satisfies

$$p = \delta(\mu_L p + \mu_H V_H),$$

since an asset owner who finds herself to have low matching quality will sell her asset and receive δp . Therefore, $p = \delta\mu_H V_H / (1 - \delta\mu_L)$. It is optimal for an asset owner with quality $(a_L + \epsilon)$ to sell her asset if

$$(1 - \delta)p \geq \max_e \left((a_L + \epsilon)e - \frac{\alpha}{2}e^2 - b \right).$$

Assume that the constraint that $e \leq 1$ is not binding. Then the optimal effort on the right hand side of the inequality is $(a_L + \epsilon)/\alpha$, and the inequality is satisfied when

$$a_L \leq \sqrt{\frac{2\alpha\delta\mu_H(1-\delta)V_H + 2\alpha(1-\delta\mu_L)b}{(1-\delta\mu_L)}} - \epsilon = \sqrt{\frac{\delta\mu_H(2a_H - \alpha)\alpha + 2(1-\delta)\alpha b}{1-\delta\mu_L}} - \epsilon = a_{LS} - \epsilon. \quad (2.7.5)$$

Assumption 2 implies that the effort $(a_L + \epsilon)/\alpha$ is indeed below one. Therefore, $q_{0+}(a_L) = 1$ when $a_L < a_{LS}$.

Next suppose that all low-quality asset owners keep their assets. Then the payoff of an asset owner with quality $a_L + \tilde{\epsilon}$ is at most $\frac{1}{1-\delta} [\max_e ((a_L + \tilde{\epsilon})e - \frac{\alpha}{2}e^2 - b)] = \frac{1}{1-\delta} \left[\frac{(a_L + \tilde{\epsilon})^2}{2\alpha} - b \right]$ where the constraint that $e \leq 1$ is ignored. Therefore, the asset price

$$p \leq \delta(\mu_H V_H + (1-\delta)^{-1}\mu_L E[(2\alpha)^{-1}(a_L + \tilde{\epsilon})^2 - b]) = \delta \left[\mu_H V_H + \frac{\mu_L}{1-\delta} \left(\frac{a_L^2}{2\alpha} - b + \frac{\epsilon^2}{6\alpha} \right) \right],$$

where $\tilde{\epsilon}$ is uniformly distributed on $[-\epsilon, \epsilon]$. It is optimal for an asset owner with quality $(a_L - \epsilon)$ to keep her asset if

$$(1 - \delta)p \leq \max_e \left((a_L - \epsilon)e - \frac{\alpha}{2}e^2 - b \right).$$

Assuming that the constraint that $e \leq 1$ is not binding, we see that the above inequality holds if

$$\frac{(a_L - \epsilon)^2}{2\alpha} - b \geq (1 - \delta)\delta\mu_H V_H + \delta\mu_L \left(\frac{a_L^2}{2\alpha} - b + \frac{\epsilon^2}{6\alpha} \right).$$

It is straight forward to show that this inequality holds when $a_L \geq a_{LS} + O(\epsilon)$, and $(a_L - \epsilon)/\alpha$ is indeed below one for $(a_L - \epsilon)$ in a neighborhood of a_{LS} under Assumption 2. Therefore, $q_{0+}(a_L) = 0$ for $a_L > a_{LS}$.

Proof of Proposition 7. Suppose that the fraction of low-quality asset owners who sell their assets is q . Let $\hat{\epsilon}(q) = 2q\epsilon - \epsilon$, the q -quantile of the uniform distribution on $[-\epsilon, \epsilon]$. Then asset owners with quality lower than $a_L + \hat{\epsilon}(q)$ will sell their assets. Therefore, the asset price is

$$p = \frac{\delta}{1 - \delta\mu_L q} \{ \mu_H V^{(w)}(a_H, p) + \mu_L (1 - q) E[V^{(w)}(a_L + \tilde{\epsilon}, p) | \tilde{\epsilon} \geq \hat{\epsilon}(q)] \}, \quad (2.7.6)$$

where $\tilde{\epsilon}$ is uniformly distributed on $[-\epsilon, \epsilon]$. Since $V^{(w)}(a_L + \tilde{\epsilon}, p) \leq V^{(w)}(a_H, p)$ for all $\tilde{\epsilon}$,

$$p \leq \frac{\delta}{1 - \delta\mu_L q} (1 - \mu_L q) V^{(w)}(a_H, p) \leq \delta V^{(w)}(a_H, p).$$

Suppose that the no-reneging constraint Eq. (2.3.6) is not binding for $a = a_H$. Then $V^{(w)}(a_H, p) = V_H = a_H - \frac{\alpha}{2} - b$, where V_H is defined in Eq. (2.7.4). It follows that $\alpha \leq \delta(V_H - \delta V_H) \leq \delta(V_H - p)$ under Assumption 2, so $V^{(w)}(a_H, p)$ is indeed equal to V_H and an employee working for a high-quality asset owner always exerts effort one.

Eq. (2.7.6) can now be rewritten as

$$p = \frac{\delta}{1 - \delta\mu_L q} [\mu_H V_H + \mu_L (1 - q) V^{(w)}(a_L + \hat{\epsilon}(q), p) + \eta_\epsilon(q)], \quad (2.7.7)$$

where $\eta_\epsilon(q) = \mu_L (1 - q) E[V^{(w)}(a_L + \tilde{\epsilon}, p) | \tilde{\epsilon} \geq \hat{\epsilon}(q)] - \mu_L (1 - q) V^{(w)}(a_L + \hat{\epsilon}(q), p)$. It is easy to show that $\eta_\epsilon(q) \geq 0$ for all q and $\bar{\eta}_\epsilon \equiv \sup_{q \in [0, 1]} \eta_\epsilon(q) \rightarrow 0$ as $\epsilon \rightarrow 0$. Let

$$s = \delta[V^{(w)}(a_L + \hat{\pi}(q), p) - p]. \quad (2.7.8)$$

Combining Eqs. (2.7.7) and (2.7.8) to eliminate p yields that

$$V^{(w)}(a_L + \hat{\pi}(q), p) = \frac{1}{\delta(1 - \delta\mu_L)} [(1 - \delta\mu_L q)s + \delta^2 \mu_H V_H + \delta^2 \eta_\epsilon(q)]. \quad (2.7.9)$$

The program Eqs. (2.3.5)-(2.3.6) at $a_L + \hat{\epsilon}(q)$ can now be rewritten as

$$\frac{1}{\delta(1-\delta\mu_L)}[(1-\delta\mu_Lq)s + \delta^2\mu_H V_H + \delta^2\eta_\epsilon(q)] = \max_{e \in [0,1]} \frac{1}{1-\delta} \left(ae - \frac{\alpha}{2}e^2 - b \right);$$

s.t. $\alpha e \leq s$,

where $a = a_L + \hat{\epsilon}(q)$. Assume that the constraint that $e \leq 1$ is not binding and the no-renegeing constraint is binding. Then

$$\frac{1}{\delta(1-\delta\mu_L)}[(1-\delta\mu_Lq)s + \delta^2\mu_H V_H + \delta^2\eta_\epsilon(q)] = \frac{1}{1-\delta} \left(\frac{as}{\alpha} - \frac{s^2}{2\alpha} - b \right).$$

This quadratic equation has a solution if and only if

$$a_L + \hat{\epsilon}(q) \geq \frac{(1-\delta)(1-\delta\mu_Lq)}{\delta(1-\delta\mu_L)}\alpha + \sqrt{a_{LS}^2 + \frac{2\alpha(1-\delta)\delta}{1-\delta\mu_L}\eta_\epsilon(q)}, \quad (2.7.10)$$

where we have used Eqs. (2.7.4) and (2.4.4) to eliminate V_H and b . Since Eqs. (2.3.5)-(2.3.6) should not have a solution for lower a , the above inequality holds with equality unless $q = 0$. When $q = 1$, Eq. (2.7.10) should not hold. Sending $\epsilon \rightarrow 0$ yields that

$$a_L = \frac{(1-\delta)(1-\delta\mu_Lq)}{\delta(1-\delta\mu_L)}\alpha + a_{LS},$$

when $q \in (0, 1)$, and a_L is weakly greater than the right hand side when $q = 0$ and a_L is weakly smaller than the right hand side when $q = 1$.

Proof of Lemma 2. The worker's incentive compatibility constraint leads to two first-order conditions:

$$\begin{aligned} c'(e) &= \cos\theta[e(m(1,1) - m(1,0)) + (1-e)(m(0,1) - m(0,0))] + \\ &\quad + (e \cos\theta + \epsilon \sin\theta)[m(1,1) - m(1,0) - m(0,1) + m(0,0)] + m(1,0) - m(0,0); \\ c'_\epsilon(\epsilon) &= \sin\theta[e(m(1,1) - m(1,0)) + (1-e)(m(0,1) - m(0,0))]. \end{aligned}$$

I use the second equation to eliminate m from the first term on the right hand side of the first equation and rewrite the system as

$$\begin{aligned} c'(e) &= c'_\epsilon(e) \cot \theta + (e \cos \theta + \epsilon \sin \theta)[m(1, 1) - m(1, 0) - m(0, 1) + m(0, 0)] + m(1, 0) - m(0, 0); \\ c'_\epsilon(\epsilon) &= \sin \theta[e(m(1, 1) - m(1, 0)) + (1 - e)(m(0, 1) - m(0, 0))]. \end{aligned}$$

For a fixed (e, ϵ) , the above is a linear system of the four components of m . I solve for $m(1, 1)$ and $m(0, 1)$ assuming that $m(1, 0)$ and $m(0, 0)$ are given. This can be done because the coefficient matrix of $(m(1, 1), m(0, 1))$ is

$$\begin{pmatrix} e \cos \theta + \epsilon \sin \theta & -e \cos \theta - \epsilon \sin \theta \\ e \sin \theta & (1 - e) \sin \theta \end{pmatrix},$$

whose determinant is always positive unless $e = \epsilon = 0$, which would imply that $V(a, p, \theta) = 0$. It is also clear from the first equation that $(m(1, 1) - m(0, 1))$ is decreasing in $m(1, 0)$, and $m(1, 1) - m(0, 1) \rightarrow \mp \infty$ as $m(1, 0) \rightarrow \pm \infty$.

Now adding a constant simultaneously to the four components of m does not change any of the constraints, so $m(0, 0)$ can be assumed to be zero. Since the map $m(1, 0) \mapsto m(1, 1) - m(0, 1)$ maps the real line onto the real line, when the no-reneging constraint is not binding, it is always possible to choose $m(1, 0)$ so that $m(1, 1) - m(0, 1) = m(1, 0)$. When the no-reneging constraint is binding, at optimality $\max\{m(1, 1) - m(0, 1), m(1, 0) - m(0, 0)\}$ must be minimized among all m that satisfies the first-order conditions. Since $(m(1, 1) - m(0, 1))$ is decreasing in $m(1, 0)$, at optimality $m(1, 1) - m(0, 1) = m(1, 0) = m(1, 0) - m(0, 0)$.

Proof of Proposition 8. The proofs of Proposition 2 and Theorem 2 go through with only one modification: at the beginning of the proof of Theorem 2 one can now show that $p_t \geq (1 - \delta)V_f(a, \theta) + \delta p_{t+1}$ in environment θ , as every asset owner can choose to use the optimal formal contract to produce for one period and then sell her asset, and using this one can prove that following fact is still true: if an asset owner decides to sell her asset at some point, it is never optimal for her to wait.

Proof of Proposition 9. Suppose that for some $a > a_c(p)$ it is optimal to use the optimal

purely formal contract. Then $V_f(a, \theta) \geq p$. In the program Eqs. (2.5.1)-(2.5.3), the worker's incentive compatibility constraint, Eq. (2.5.2), can be replaced by the following first-order conditions:

$$\begin{aligned} c'(e) &= m_f \cos \theta + m_r; \\ c'_\epsilon(\epsilon) &= m_f \sin \theta. \end{aligned}$$

Therefore, $m_r = c'(e) - c'_\epsilon(\epsilon) \cot \theta$. Since the entrepreneur's payoff from a stationary relational contract is $(1 - \delta)^{-1}[ae - c(e) - c_\epsilon(\epsilon) - b]$, there exists a relational contract for which the no-renegeing constraint Eq. (2.5.3) is not binding if and only if there exists (e, ϵ) such that

$$c'(e) - c'_\epsilon(\epsilon) \cot \theta < \delta[(1 - \delta)^{-1}(ae - c(e) - c_\epsilon(\epsilon) - b)] - \delta V_f(a, \theta).$$

This is equivalent to the condition that $\gamma_\theta(a) > (1 - \delta)^{-1}b + V_f(a, \theta)$, where

$$\gamma_\theta(a) = \max_{e, \epsilon} \frac{1}{1 - \delta} [ae - c(e) - c_\epsilon(\epsilon)] - \delta^{-1} [c'(e) - c'_\epsilon(\epsilon) \cot \theta]. \quad (2.7.11)$$

However, the definition of V_f directly implies that $(1 - \delta)^{-1} + V_f(a, \theta)$ is the value of the above program with the additional constraint that $c'(e) - c'_\epsilon(\epsilon) \cot \theta = 0$, so $\gamma_\theta(a) \geq (1 - \delta)^{-1} + V_f(a, \theta)$ for all a , and strictly so for all but one a since the optimal ϵ in Eq. (2.7.11) is independent of a , and thus the equality $c'(e) - c'_\epsilon(\epsilon) \cot \theta = 0$ holds at optimality for at most one a . As long as $\gamma_\theta(a) > (1 - \delta)^{-1}b + V_f(a, \theta)$, there exists a relational contract in which the no-renegeing constraint Eq. (2.5.3) is not binding, which implies that the value of this relational contract is strictly above $V_f(a, \theta)$; therefore, the value of the optimal relational contract is also strictly above $V_f(a, \theta)$.

Chapter 3

Optimal Mechanisms with Information Acquisition

3.1 Introduction

Contract theory and organizational economics have been concerned with asymmetric information for decades, and this literature has been fruitful in identifying inefficiencies in organizations and transactions, as well as drawing implications about organizational design. However, for simplicity, most papers assume that agents are *endowed* with different information; for example, one party has superior information about the state of nature at the beginning of the game. This paper explores the situation in which one party may gather some information, but only at a cost.

In many real world applications, asymmetric information arises from differential cost of gathering and processing information. For example, in the decision making literature, dating back to Crawford and Sobel (1982), it is often assumed that the agent has superior information about the state of nature. Yet in the real world, it is not impossible for the principal to gather the same information, but the principal chooses not to do either because she lacks the expertise to process the information effectively, or her opportunity cost of time is bigger than the agent's. Therefore, assuming that the agent is endowed with superior information is only a simplification. In fact, in many papers, the agent is referred to as an "expert" to emphasize that he has the specific skill to gather and process information, but he is not an "insider" in the sense that he only has access to the public information, which should also be accessible to the principal. In all these

settings, the real difference between the agent and the principal is the cost of acquiring information, instead of access to information.

On the other hand, even though the agent can acquire information at a lower cost than the principal, he often needs to spend time and resources to do so. In the real world, it is not uncommon for a consultant to spend months collecting data before he can provide meaningful advice. The same is true in the classical regulation literature, dating back to Baron and Myerson (1982) and Laffont and Tirole (1986). The product to be produced by the firm is often a specialized one, so it is not obvious that the firm would have superior information about the cost of production without actually producing it. Consequently, the firm may need to perform some cost analysis before signing the contract in order to learn the cost of production, which is costly. Therefore, it would be interesting to check whether the result provided in the literature of information transmission, expert advice and regulation is robust when the agent's superior information comes at a cost, and the agent's effort to acquire information is not contractible.

This paper considers the classical screening problem in which the principal offers the agent a menu of allocations and transfers, except that the agent can choose his effort to acquire information. The effort of information acquisition, as well as the information itself, is assumed to be uncontractible. The goal is to characterize the optimal contract (mechanism) from the principal's perspective. In some simple cases, a very detailed characterization can be obtained and contrasted with the situation when the agent's superior information is costless.

Under suitable risk-neutrality and single-crossing assumptions on the agent's payoff function, the paper shows that the agent is always "information-loving": he would like to acquire the maximum amount of information if information were free. The reason is that under these assumptions the agent's payoff in the mechanism is a convex function of his posterior expectation of the signal. The result mentioned above then becomes a corollary of Jensen's inequality. Furthermore, the agent's decision on information acquisition has a close analogue in financial economics: the agent's signal acts as the value of an underlying asset and his payoff in the mechanism is the exercise payoff of an options contract. Therefore, the agent's information acquisition problem is reduced to the problem of pricing an options contract.

The characterization of the optimal mechanism relies on the solution to the "options pricing" problem mentioned above, but the characterization is simple and gives many economic insights when the information acquisition is assumed to be binary: the agent acquires either no information or all information. In this case, the paper shows that the optimal mechanism depends on the cost of information acquisition. When the cost is low, with costless information (often referred to as the "second best" in the screening literature) remains implementable. When the cost is higher, deviation from the second best is required to give the agent more informational rent. In particular, the allocation to the highest type may no longer be efficient; instead, the allocation for higher types will be upward distorted for high information acquisition cost.

However, the deviation from the second best does not necessarily reduce the total surplus. In fact, the paper illustrates that in some cases the allocation in the optimal mechanism with costly information acquisition is more efficient than the second best allocation. This raises the following issue: when the information acquisition action is contractible, whether contracting on that is beneficial. It is shown that it contracting on information acquisition does not have to improve the total surplus.

Another situation in which characterization of the optimal mechanism is relatively easy is when the agent chooses the probability that the true distribution is revealed to him. This setting is similar to those considered in Aghion and Tirole (1997) and Zermeno (2011). In this case, the agent chooses his information acquisition action from a continuum, and it is possible to compare information generated under first best and information generated in the optimal mechanism. A sufficient condition for the latter is smaller is given in the paper.

The paper is related to the previous works on mechanism design with information acquisition. In particular, the set up in Crémer *et al* (1998b) is a specialization of the model considered in this paper. It is argued that the shadow cost of information acquisition for the principal must be smaller than 1 (Lemma 2), which is true in their modified Baron-Myerson model, but does not necessarily hold in the more general case, as illustrated in Section 3 of this paper. Crémer *et al* (1998a) and Crémer and Khalil (1992) consider strategic information acquisition prior to contracting, in the sense that the information becomes costless after the agent signs the contract with the principal. In their terminology, this paper considers "productive" information acquisition.

Shi (2008) considers the design of optimal auctions, and Lewis and Sappington (1997) considers a model with moral hazard, both of which consider the problem of endogenous information acquisition, but in settings slightly different from this paper. Bergemann and Välimäki considers efficient mechanism design, while this paper considers optimal mechanism design.

The paper is also related to the literature on decision making with information acquisition, dating back to Lambert (1986) and Demski and Sappington (1987). A common feature in that literature is that distortion in decisions may arise in the optimal contract to induce information acquisition, which is similar to the spirit of the current paper. However, in the risk neutral settings such as Inderst and Klein (2007) and Zermeno (2011), the decision maker never takes the safe action too often. This paper shows that this needs not be true in more general environments, especially when the agent also has a stake in the outcome. (See Section 3.2 for details.)

The rest of the paper is organized as follows. Section 2 presents the general setup of the model and discusses assumptions. Section 3 analyzes the two-type environment with binary information acquisition. Section 4 provides some characterization of the optimal mechanism in the more general setting. Section 5 studies two relatively tractable examples. Section 6 concludes.

3.2 The Model

In the model there is a principal (she) and an agent (he) trying to make a decision x (also referred to as the allocation) based on the state of nature θ . The payoffs of both parties depend can depend on both x and θ . It is assumed that x is the only contractible variable, and the parties have transferable utility. The principal and the agent have a common prior on θ , but the agent can acquire more information at his own cost. It is assumed that the agent acquires information after the contract is offered and before he decides whether to sign the contract. Formally, the timing is as follows:

1. The principal offers a contract.
2. The agent observes the offer, chooses an information acquisition action n from a feasible set \mathcal{N} at cost $c(n)$, and observes the signal produced by n . Based on that,

he decides whether to sign the contract. If he rejects the contract, the game ends and each party gets outside value 0.

3. If the agent chooses to contract with the principal, he chooses an allocation-transfer pair from the menu offered by the principal, which is then implemented. The agent's payoff is $u(\theta, x) + t - c(n)$ and the principal's payoff is $\pi(\theta, x) - t$.

The timing is identical to the timing in Crémer *et al* (1998b). This is the closest analogue to the classical screening problem: the agent can decide whether to sign the contract after he acquires private information, so the contract offered by the principal essentially has to satisfy the interim participation constraint. A remaining degree of freedom is whether the agent acquires information before or after the contract is offered. The difference is discussed in Crémer *et al* (1998b), and the key point is that in equilibrium the agent may choose some mixed information acquisition strategy before the contract is offered, which will make the analysis of optimal contract much harder. Therefore, in this paper, I focus on the setting in which the agent acquires information after the contract is offered. The solution concept is sequential equilibrium.

A different but related timing is that the agent and the principal have already signed a contract, but the agent is protected by limited liability, and the principal can still commit to a contract before the agent collects information. The analysis below still holds, except that the interim participation constraint will be replaced with the limited liability constraint. Limited liability is usually assumed in the decision making literature, such as Inderst and Klein (2007) and Zermeno (2011).

It is assumed that payoffs of both the agent and the principal can depend on both the decision x and the state of nature θ . On one hand, this is different from the regulation problem as in Baron and Myerson (1982), where the regulator's payoff does not depend on the state of nature (the agent's cost of production). This is not an innovation of this paper, as the procurement model considered in Lewis and Sappington (1997), among other models, also allows the principal's payoff to depend on the state of nature. Unlike Crémer *et al* (1998b) which considers a variant of the Baron-Myerson model, the shadow cost of information acquisition for the principal can be bigger than one, which makes efficiency at the top fail. On the other hand, this is different from the decision making literature dating back to Lambert (1986), where the agent's payoff does not depend on

the state of nature (the probability that each outcome occurs). In fact, in the decision making literature, it is usually assumed that the agent's payoff does not depend on the decision x either. As shown later in this paper, the general setting delivers a richer set of phenomena. Unlike that literature, there can be both underinvestment and overinvestment in the risky project.

The set \mathcal{N} of information acquisition actions is assumed to have the following structure. There is a stochastic process $\{Z_n\}_{n \in N}$ in discrete or continuous time. The probability distribution of this stochastic process depends on the state of nature, θ . The set of actions \mathcal{N} is a set of stopping times of this stochastic process. If the agent chooses $n \in \mathcal{N}$, then he observes $\{z_m\}_{m \leq n}$. For example, \mathcal{N} can be N itself, in which case the agent has to decide when to stop acquiring information before he observes anything; \mathcal{N} can also be the set of all stopping times of the stochastic process, in which case the agent can decide when to stop contingent on the signals already observed. Denote the filtration generated by this stochastic process by $\{\mathcal{F}_n\}_{n \in N}$.

The seemingly strange assumption about \mathcal{N} is designed to capture the idea that the agent never "forgets": if $n < n'$, then $\mathcal{F}_n \subset \mathcal{F}_{n'}$, and the agent has more information under n' than under n . The simplest example is that $\mathcal{N} = N = \{0, 1\}$, $z_0 = 0$ and $z_1 = \theta$. In this case the agent chooses either not to acquire information ($n = 0$) or to acquire all the relevant information ($n = 1$). This is the setting used in both Crémer *et al* (1998b) and in Lambert (1986). This setting is tractable and generates clear economic insights, so a substantial fraction of this paper will focus on this setting too. However, some of the general characterization applies to more general information process and the setting used here leads to a natural analogue to the option pricing problem in financial economics, as presented in Section 4. Here are two examples of the information process:

Example 1. Let $\Theta \subset \Delta(\Xi)$ where Ξ is a finite set, $u(\theta, x) = \sum_{\xi \in \Xi} \tilde{u}(\xi, x)\theta(\xi)$ and $\pi(\theta, x, t) = \sum_{\xi \in \Xi} \tilde{\pi}(\xi, x, t)\theta(\xi)$ for some functions \tilde{u} and $\tilde{\pi}$. Let the agent and the principal has the Dirichlet prior about θ , and z_n a random draw from Ξ with probability distribution θ . Then the posterior belief about θ is still a Dirichlet distribution.

Example 2. Let $N = [0, 1]$ and τ be a random variable independent of θ and uniformly distributed on $[0, 1]$. Assume that $z_n = z_0 \notin \Theta$ for $n < \tau$ and $z_n = \theta$ for $n \geq \tau$. In this case, the agent chooses the probability that the state of nature is observed. This is the setting used in Aghion and Tirole (1997) and also discussed in the context of decision making in Zermeno (2011).

3.3 The Two-Type Case

3.3.1 Additional distortion and inefficiency at the top

This section considers the case when $\Theta = \{\theta_L, \theta_H\}$ where $\theta_L < \theta_H$. Assume that $X \subset \mathbb{R}$ and that $u(\theta, x)$ is increasing in θ and has increasing differences in θ and x . Consider the particular information structure in which $N = \{0, 1\}$, $z_0 = 0$ and $z_1 = \theta$. In other words, the agent can choose to observe nothing or observe the state of the world θ . The common prior is that $\Pr(\theta = \theta_H) = q$. Let $K = c(1) - c(0)$, the cost of observing θ . I call a mechanism *informative* if it induces the agent to choose $n = 1$, and call a mechanism *uninformative* otherwise. By the revelation principal, an informative mechanism can be written as $\{(x_L, t_L), (x_H, t_H)\}$. The next proposition characterizes all incentive compatible informative mechanisms.

Proposition 21. A mechanism $\{(x_L, t_L), (x_H, t_H)\}$ is incentive compatible if and only if both of the following constraints are satisfied:

$$u(\theta_L, x_L) + t_L \geq 0; \quad (3.3.1)$$

$$u(\theta_H, x_L) - u(\theta_H, x_H) + \frac{K}{q} \leq t_H - t_L \leq u(\theta_L, x_L) - u(\theta_L, x_H) - \frac{K}{1-q}. \quad (3.3.2)$$

Proof. The standard IC constraints read

$$u(\theta_L, x_L) + t_L \geq u(\theta_L, x_H) + t_H;$$

$$u(\theta_H, x_H) + t_H \geq u(\theta_H, x_L) + t_L.$$

These imply that the participation constraint for θ_H is redundant, so the only relevant participation constraint is Eq. (3.3.1). The IC constraints are equivalent to the following inequality:

$$u(\theta_H, x_L) - u(\theta_H, x_H) \leq t_H - t_L \leq u(\theta_L, x_L) - u(\theta_L, x_H). \quad (3.3.3)$$

The agent has incentive to acquire information if and only if

$$q[u(\theta_H, x_H) + t_H] + (1-q)[u(\theta_L, x_L) + t_L] \geq \max_{i \in \{L, H\}} \{q[u(\theta_H, x_i) + t_i] + (1-q)[u(\theta_L, x_i) + t_i]\}.$$

This constraint is equivalent to Eq. (3.3.2) and implies Eq. (3.3.3). \square

Several useful observations can be made:

1. In the mechanism design problem without information acquisition, $t_H - t_L = u(\theta_H, x_L) - u(\theta_H, x_H)$ at optimality ("second best"). When information is costly, this scheme fails to induce information acquisition no matter how small the cost K is. The intuition is that in that scheme the type θ_H is indifferent between telling the truth and reporting θ_L , and therefore there is no gains from always reporting θ_L to telling the truth, which implies that the agent has no incentive to acquire information. Therefore, an informative mechanism must concede more rent to the high type.
2. Proposition 1 implies that an allocation rule (x_L, x_H) can be implemented by some transfer rule if and only if

$$u(\theta_H, x_L) - u(\theta_H, x_H) + \frac{K}{q} \leq u(\theta_L, x_L) - u(\theta_L, x_H) - \frac{K}{1-q}. \quad (3.3.4)$$

When $K = 0$, this is equivalent to the monotonicity condition $x_L \leq x_H$ given the single-crossing assumption I have made. When $K > 0$, Eq. (3.3.4) imposes a stronger condition than monotonicity. Intuitively, the allocations for different types have to be sufficiently different to justify the cost of acquiring information.

3. The additional distortion created by information acquisition depends on K/q and $K/(1-q)$. Therefore, the distortion is bigger when q is closer to 0 or 1. In fact, when q is very close to 0 or 1, it would be optimal not to acquire information because the distortion caused by information acquisition incentive is too big and the benefit of screening is small.

Now consider the optimal informative mechanism. Clearly, the principal chooses t_L and t_H such that Eq. (3.3.1) and the left half of Eq. (3.3.2) holds with equality. Therefore, the optimal allocation rule solves the following program (P):

$$\begin{aligned} \max_{x_L, x_H} \quad & q[\pi(\theta_H, x_H) + u(\theta_L, x_L) + u(\theta_H, x_H) - u(\theta_H, x_L)] + (1-q)[\pi(\theta_L, x_L) + u(\theta_L, x_L)] - K \\ \text{s.t.} \quad & u(\theta_L, x_L) - u(\theta_L, x_H) - \frac{K}{1-q} \geq u(\theta_H, x_L) - u(\theta_H, x_H) + \frac{K}{q}. \end{aligned}$$

Let (x_L^*, x_H^*) be the solution to the program without the constraint; i.e. the second best allocation in the standard screening problem. Let $\Pi(K)$ be the principal's payoff in the optimal informative mechanism.

Proposition 22. (x_L^*, x_H^*) is the allocation in the optimal informative mechanism if and only if

$$\frac{K}{q} + \frac{K}{1-q} \leq u(\theta_L, x_L^*) + u(\theta_H, x_H^*) - u(\theta_L, x_H^*) - u(\theta_H, x_L^*). \quad (3.3.5)$$

In this case $\Pi'(K) = -1$. Assume in addition that X is an interval, $\pi(\theta, \cdot), u(\theta, \cdot)$ are continuously differentiable on X and $\pi(\theta, \cdot) + u(\theta, \cdot)$ is concave, then in the optimal informative mechanism, $x_H \geq x_H^*$ and $x_L \leq x_L^*$ when Eq. (3.3.5) is violated, in which case $\Pi'(K) < -1$.

Proof. The first part is obvious from Proposition 1. Let $S(\theta, x) = \pi(\theta, x) + u(\theta, x)$. Then the first-order condition for x_H can be written as

$$qS_x(\theta_H, x_H) + \lambda[u_x(\theta_H, x_H) - u_x(\theta_L, x_H)] = 0,$$

where λ is the Lagrange multiplier, which is positive when Eq. (3.3.5) is violated. By the single-crossing property of u , $u_x(\theta_H, x_H) - u_x(\theta_L, x_H) \geq 0$, so $S_x(\theta_H, x_H) \leq 0$. Since $S_x(\theta_H, x_H^*) = 0$ and $S(\theta_H, \cdot)$ is concave, $x_H \geq x_H^*$. Similarly, x_L satisfies the following first-order condition:

$$(1-q)S_x(\theta_L, x_L) - (\lambda+q)[u_x(\theta_H, x_L) - u_x(\theta_L, x_L)] = 0,$$

and x_L^* satisfies the same condition with $\lambda = 0$. Since $u_x(\theta_H, x_L) - u_x(\theta_L, x_L) \geq 0$, $x_L \leq x_L^*$.

To obtain $\Pi'(K)$, apply the envelop theorem to the program (P):

$$\Pi'(K) = -1 - [q(1-q)]^{-1}\lambda,$$

where $\lambda \geq 0$ is the Lagrange multiplier of the constraint. Therefore, $\Pi'(K) = -1$ when the constraint is not binding and $\Pi'(K) < -1$ when the constraint is binding. \square

Proposition 2 implies that when K is below some critical value, introducing the cost of acquiring information does not change the optimal allocation. The only distortion is that

the high type receives a higher rent. However, when K is bigger, the optimal allocation will be different, and in particular, the "efficiency at the top" fails: the allocation to the high type receives upward distortion. Furthermore, since an increase in K not only increases the agent's rent, but also creates more distortion in the allocation, the shadow cost of information acquisition for the principal, $-\Pi(K)$, is bigger than one when the second best allocation is not implementable.

A similar result is delivered in Lewis and Sappington (1997) in the presence of moral hazard, but the intuition remains the same: when the cost of information is big, to induce information acquisition the contract has to make it costly enough for the low type to imitate the high type, which leads to upward distortion for the high type and (further) downward distortion for the low type. However, the same result does not arise in the Baron-Myerson like model considered in Crémer *et al* (1998b) as in their setting the shadow cost of information acquisition cannot exceed 1 (their Lemma 2). In this sense, the phenomena that can arise in the general setting considered in this paper is richer than in Crémer *et al* (1998b). The caveat is that when K is big, it might not be optimal for the principal to induce information acquisition in the first place, so the "inefficiency at the top" never arises. The next subsection illustrates with an example that the inefficiency at the top can arise even when it is optimal to induce information acquisition.

3.3.2 Investing in a risky project: an example

In this example, the principal is the general manager of an organization, and the agent is the head of a division. The division head proposes a project, which would yield total surplus xR if it succeeds and 0 if it fails, where $x \geq 0$ is investment in this project. The relevant information is the probability θ with which the project will succeed. Due to the lack of expertise, the manager has to delegate the acquisition of information to the division head, but the investment can only be funded by the general manager, at a convex cost $I(x)$. The intuition is that the manager may have some alternative use of the fund, and the marginal opportunity cost of fund is increasing in the investment. The division head has share α of the total surplus, and the manager has share $(1 - \alpha)$, where $\alpha \in (0, 1)$ is exogenously given. To sum up, the agent's payoff is $u(\theta, x) + t = \alpha\theta xR + t$, and the principal's payoff is $\pi(\theta, x) - t = (1 - \alpha)\theta xR - I(x) - t$. The investment x is

contractible, but the outcome of the project is not, so transfer can only depend on x , but not on θ . The reason that the outcome is not contractible might be that the outcome is hard to evaluate. For example, the project might be some organizational innovation, whose benefit can only be evaluated against some counterfactual, which is hard to put into a contract. Another possible interpretation is that there is some long-run benefit of the project that cannot be verified by the time the transfer has to be paid.

Assume that by investing $K > 0$, the agent can learn whether the probability of success is θ_H or θ_L , where $\theta_H > \theta_L$. The common prior is that $\Pr(\theta = \theta_H) = q$. Let $\theta_0 = q\theta_H + (1 - q)\theta_L$. Let $g(b) = \operatorname{argmax}_{x \geq 0} bx - I(x)$. It is easy to see that g is increasing in b . When I is continuously differentiable and strictly convex, g is simply the inverse function of I' . Then the first best investment is given by $x^{FB}(\theta) = g(\theta R)$, if both the information θ were costless and contractible. The second best investment is given by

$$\begin{aligned} x_H^{SB} &= g(\theta_H R); \\ x_L^{SB} &= g\left(\theta_L R - \frac{\alpha q}{1 - q}(\theta_H - \theta_L)R\right). \end{aligned}$$

Not surprisingly, there is no distortion for the high type, and downward distortion for the low type. An immediate observation is that the loss from adverse selection is increasing in α . Therefore, the organization is more efficient when α is smaller, from the perspective of *ex ante* organizational design. If α is the attribute of the project instead of the attribute of the organization, then the comparative statics imply that projects whose benefit accrues more to the principal are handled more efficiently in the organization. This will no longer hold when the information is costly.

When the information is costly and not contractible, the optimal informative and uninformative mechanism are characterized by the following proposition:

Proposition 23. When $K \leq K_c$, the allocation in the optimal informative mechanism is (x_H^{SB}, x_L^{SB}) ; when $K > K_c$, $x_H \geq x_H^{SB}$ and $x_L^{SB} \leq x_L^{SB}$ in the optimal informative mechanism, where

$$K_c = q(1 - q)\alpha R(\theta_H - \theta_L)(x_H^{SB} - x_L^{SB}). \quad (3.3.6)$$

In the optimal uninformative mechanism, the optimal investment is

$$x = \begin{cases} g(\theta_0 R), & \text{if } K \geq (1 - q)\alpha R(\theta_0 - \theta_L)g(\theta_0 R); \\ g((1 - \alpha)\theta_0 R + \alpha\theta_L R), & \text{otherwise.} \end{cases}$$

Proof. The characterization of the optimal informative mechanism follows directly from Proposition 2. The optimal uninformative contract (x, t) solves

$$\begin{aligned} \max_{x,t} \quad & (1 - \alpha)\theta R x - I(x) - t; \\ \text{s.t.} \quad & \alpha\theta_0 R x + t \geq 0; \\ & \max\{0, -t - \alpha\theta_L R x\}(1 - q) \leq K. \end{aligned}$$

The second constraint says that the agent has no incentive to acquire information: the benefit of acquiring information is that he avoids the loss $\max\{0, -(\alpha\theta_L R x + t)\}$ when $\theta = \theta_L$ (by rejecting the contract when $\theta = \theta_L$), and the cost is K . Solving the program gives the equations in the proposition. \square

The characterization of the optimal uninformative mechanism is the analogue of the corresponding result in Crémer *et al* (1998b). For big K , the investment maximizes the total surplus for $\theta = \theta_0$, the prior belief. For smaller K , in order to "deter" information acquisition, investment is reduced to curb the potential loss that the agent will incur when his type is low, which is equal to his benefit of acquiring information.

Since the agent's payoff is proportional to αR , it should not be surprising that K_c is proportional to αR also. This implies that sometimes the principal's payoff is higher for higher α , as when K is close to K_c , a higher α helps the principal to avoid additional distortion from information inducing. (When X is a discrete set, the principal's loss may be substantial when K exceeds K_c .)

Note that x_H^{SB} is increasing in θ_H and x_L^{SB} is increasing in θ_H and decreasing in θ_L , so Eq. (3.3.6) implies that K_c is increasing in θ_H and decreasing in θ_L . This is intuitive: when θ_H is bigger and θ_L is smaller, the informational rent for the high type is bigger, and the cost for the uninformative agent to immitate the high type is also bigger, so the total benefit of acquiring information is bigger, and can thus sustain a higher cost of information. Finally, since x_L^{SB} is bounded from below and x_H^{SB} is independent of q ,

K_c approaches zero if q approaches zero or one. This is also intuitive: when the agent is almost sure about his type *ex ante*, he does not want to acquire information under the second-best contract.

To see that upward distortion can arise in the optimal mechanism, consider a continuously differentiable and strictly convex cost function I for which $I'(0) \in (\theta_0 R, \theta_H R)$. Then $g(\theta_0 R) = 0$, and no investment is made in the optimal uninformative mechanism and the principal's payoff is zero. The principal's payoff from the optimal informative mechanism is $q\theta_H R g(\theta_H R) - I(g(\theta_H R)) - K$. By Eq. (3.3.6), $K_c = q(1 - q)\alpha R(\theta_H - \theta_L)g(\theta_H R)$, which is proportional to α , so the principal's payoff is positive when $K = K_c$ if α is sufficiently small. When $K > K_c$, in the optimal informative mechanism $x_H = K/[q(1 - q)\alpha R(\theta_H - \theta_L)]$ since Eq. (3.3.4) must hold with equality and Proposition 3 implies that $x_L = 0$. Since I is continuous in x_H and x_H is continuous in K , the principal's payoff is also continuous in K . Therefore, there exists an $\epsilon > 0$ such that for $K \in (K_c, K_c + \epsilon)$, the principal's payoff is positive and $x_H > x_H^{SB}$ in the optimal informative mechanism.

Proposition 3 implies that when $K > K_c$ and it is optimal for the principal to induce information acquisition, there is over-investment when $\theta = \theta_H$ and underinvestment in $\theta = \theta_L$. In other words, the principal invests too much when the signal is favorable and too little when the signal is unfavorable, compared to both second best and first best. The result partially extends to the case with general type space, as Section 5 shows. This is in contrast to the result obtained in the decision making literature such as Inderst and Klein (2007) and Zermeno (2011). There the intuition is that the principal can find out the true state of nature only if she takes the risky action, so to motivate information acquisition, she takes the risky action too often. In the investment problem discussed here, both over-investment and under-investment in the risky project can happen, even though it is still true that the principal can find out the true state of nature (whether the project succeeds) only if she invests in the risky project.

From the mechanism design perspective, it is not surprising that under-investment, or downward distortion, can happen under asymmetric information. The reason why in the decision making literature taking the safe action too often cannot happen is that in some sense there is no room for downward distortion. To fit the decision making problem into the framework here, I extend the allocation space to $\{0, 1\} \times \mathbb{R}$, where the

first component d is the decision (0 for the safe action and 1 for the risky action), and the second component x is the bonus paid to the agent if the risky action is taken and successful. As usually done in the decision making literature, I assume that the agent does not derive utility from the decision, so his utility is given by $\theta x 1_{d=1} + t$ and $d = 0$ is equivalent to $x = 0$ for him. Under the first best (when the agent's superior information becomes contractible), $x = 0$ always no matter what d is. Therefore, changing d from 1 to 0 is not a distortion that changes the agent's incentive. Consequently, if any distortion needs to be created to satisfy the agent's IC constraints, it should be making x nonzero and changing d from 0 to 1 when necessary. Usually $x < 0$ is ruled out by model assumptions, so the only possible distortion is to pay positive bonus to the agent and sometimes takes $d = 1$ even when the safe action yields higher total surplus in order to give the agent more bonus. (The decision making literature usually assumes limited liability instead of participation constraint, so the framework here needs to be modified slightly to adapt to that scenario.)

3.4 The Linear Model

Unfortunately, the mechanism design problem with more than two types is harder to solve. Therefore, I focus on the special case when the agent's payoff is linear in his type, as it is easier to obtain some useful results and insights in this case. Specifically, I assume that $\Theta, X \subset \mathbb{R}$ and $u(\theta, x) = \theta h_1(x) + h_0(x)$, where h_1 is a nonnegative and increasing function. Assume that there exists $x_0 \in X$ such that $h_1(x_0) = h_0(x_0) = 0$ so that letting the agent leave is a possible allocation. It is straight forward to extend the analysis to the case where $u(\theta, x) = g(\theta)h_1(x) + h_0(x)$ for a positive and strictly increasing g , as one can refer to $g(\theta)$ as the "state of nature".

3.4.1 Value of information and the role of the prior

In general, $E_n[\theta]$ can be any convex combination of elements of Θ , so it is necessary to consider all θ in the convex hull of Θ . To simplify the notation, I denote the convex hull of Θ by Θ in this section. (Equivalently, I assume that Θ is an interval.)

The agent's conditional expected payoff is given by

$$E_n[u(\theta, x) + t] = E_n[\theta]h_1(x) + h_0(x) + t.$$

Therefore, if the agent chooses to stop acquiring information at time n , he acts as if the true state of the world were $E_n[\theta]$. By the revelation principle, the principal has to offer $\{(x(\theta), t(\theta)) : \theta \in \Theta\}$ such that

$$\theta \in \operatorname{argmax}_{\tilde{\theta}} \theta h_1(x(\tilde{\theta})) + h_0(x(\tilde{\theta})) + t(\tilde{\theta}).$$

(This is without loss of generality even if some $\theta \in \Theta$ appears with zero probability as the conditional expectation in equilibrium.)

Let

$$V(\theta) = \theta h_1(x(\theta)) + h_0(x(\theta)) + t(\theta). \quad (3.4.1)$$

The single-crossing property (in this context, the monotonicity of h_1) and the envelop theorem implies that

1. $x(\theta)$ is weakly increasing in θ ;
2. $V(\theta) = \int_{\theta_0}^{\theta} h_1(x(\tilde{\theta}))d\tilde{\theta} + V(\theta_0)$ for all $\theta_0, \theta \in \Theta$.

Combining the above, one obtains that

$$V'(\theta) = h_1(x(\theta)), \quad (3.4.2)$$

which is weakly increasing in θ . Therefore, V is a convex function.

If the agent chooses to stop observing the signal at n , his payoff (gross of the cost of acquiring information) is $V(E[\theta|\mathcal{F}_n])$. (Depending on whether the agent is allowed to choose when to stop contingent on the observed signals, n and n' may be deterministic times or stopping times of the signal process $\{z_n\}$.) Now consider two times n and n' with $n' < n$. Then Jensen's inequality implies that

$$E[V(E[\theta|\mathcal{F}_n])|\mathcal{F}_{n'}] \geq V(E[E[\theta|\mathcal{F}_n]|\mathcal{F}_{n'}]) = V(E[\theta|\mathcal{F}_{n'}]).$$

Therefore, at any point n' , the agent is always willing to acquire more information if the information were free. At time n' , the expected value of the information arriving from n' to n is given by the difference

$$E[V(E[\theta|\mathcal{F}_n])|\mathcal{F}_{n'}] - V(E[\theta|\mathcal{F}_{n'}]).$$

To simplify the notation, denote $E[\theta|\mathcal{F}_n]$ by θ_n . Then clearly $\{\theta_n\}$ is a martingale. Let \mathcal{N} be the set of all information acquisition strategies. For example, when the agent cannot decide whether to stop acquiring information contingent on the observed signals, $\mathcal{N} = N$; when the agent can decide whether to stop acquiring information in any contingent way, \mathcal{N} is the set of all stopping times of the process $\{z_n\}$. The agent's problem is

$$\max_{n \in \mathcal{N}} E[V(\theta_n) - c(n)].$$

This problem has an analogue in financial economics. Imagine that θ_n is the value of some underlying asset, and $V(\theta)$ is the payoff from exercising an option when the value of the underlying asset is θ . When the agent's has to decide the time to stop *ex ante*, the objective becomes $E[V(\theta_n)] - c(n)$ where the first term is the value of the *European* option which matures at n , so the agent is choosing among the European options mature at different time. When the agent's can decide when to stop contingent on observed signals, the agent's problem is equivalent to the problem of exercising an *American* option, and $c(n)$ plays the role of dividend in the option pricing problem.

A trivial observation is that the value of an option depends on the initial value of the underlying asset. Therefore, the agent's information acquisition strategy depends on $E[\theta|\mathcal{F}_0]$. Similarly, the agent's information acquisition strategy under a fixed allocation rule depends on his prior. Unlike the classical screening problem without information acquisition, the prior distribution of θ enters the problem in two ways. First, it affects the principal's objective function as in the classical screening problem. Secondly, it affects the agent's information acquisition strategy, and thus the design of the mechanism, which should induce information acquisition. When the principal and the agent have different priors (both of which are common knowledge), the principal needs to use the agent's prior to check whether the allocation induces an appropriate level of information

acquisition. In contrast, the agent's prior plays no role when information comes at no cost to the agent.

3.4.2 Characterizing the conditional optimal allocations

Given an information acquisition strategy $n \in \mathcal{N}$ and an allocation rule $x : \Theta \rightarrow X$, Each information acquisition strategy $n' \in \mathcal{N}$ leads to a distribution over $\Theta = [\underline{\theta}, \bar{\theta}]$, whose cdf is denoted by $F^{(n')}$. By the envelop theorem, the agent's payoff from n' is

$$\int V(\theta) dF^{(n')}(\theta) - c(n') = \int [1 - F^{(n')}(\theta)] h_1(x(\theta)) d\theta - c(n').$$

Lemma 12. The allocation rule x and the information acquisition strategy $n \in \mathcal{N}$ can be implemented if and only if $x(\theta)$ is nondecreasing in θ and

$$\int F^{(n)}(\theta) h_1(x(\theta)) d\theta + c(n) \leq \int F^{(n')}(\theta) h_1(x(\theta)) d\theta + c(n') \text{ for all } n' \in \mathcal{N}. \quad (3.4.3)$$

Proof. The standard theory of mechanism design implies that the IC constraints are equivalent to envelop theorem formula $V(\theta) = \int h_1(x(\theta)) d\theta$ combined with the monotonicity constraint. The IR constraint is binding only for type $\underline{\theta}$ since $h_1(x) \geq 0$ for all x . The inequality in the lemma is a mere restatement of the optimality of n as shown in the discussion preceding the lemma. \square

Under the information strategy n , the total transfer to the agent is

$$\int [V(\theta) - u(\theta, x(\theta))] dF^{(n)}(\theta) = - \int u(\theta, x(\theta)) dF^{(n)}(\theta) + \int [1 - F^{(n)}(\theta)] h_1(x(\theta)) d\theta.$$

Consider the principal's problem of conditional optimal allocation, defined as the optimal allocation under a given $n \in \mathcal{N}$. By Lemma 1, this allocation solves the following program:

$$\begin{aligned} \max_x \quad & \int [\pi(\theta, x(\theta)) + u(\theta, x(\theta))] dF^{(n)}(\theta) - \int [1 - F^{(n)}(\theta)] h_1(x(\theta)) d\theta, \\ \text{s.t.} \quad & x'(\theta) \geq 0; \\ & \int F^{(n)}(\theta) h_1(x(\theta)) d\theta + c(n) \leq \int F^{(n')}(\theta) h_1(x(\theta)) d\theta + c(n'), \text{ for all } n' \in \mathcal{N}. \end{aligned}$$

Solving the above program when \mathcal{N} is an infinite set, especially a continuum, may require the variational method. To get some insight, consider the case when \mathcal{N} is a finite set. Let $f^{(n)}$ be the p.d.f. associated with $F^{(n)}$. Ignoring the monotonicity constraint, one can write down the Lagrangian:

$$\mathcal{L}(x) = \int \left[S(\theta, x(\theta)) - \frac{1 - F^{(n)}(\theta)}{f^{(n)}(\theta)} h_1(x(\theta)) - \sum_{n' \in \mathcal{N}} \lambda_{n'} \frac{F^{(n)}(\theta) - F^{(n')}(\theta)}{f^{(n)}(\theta)} h_1(x(\theta)) \right] f^{(n)}(\theta) d\theta - \sum_{n' \in \mathcal{N}} \lambda_{n'} [c(n) - c(n')] \quad (3.4.4)$$

where $S(\theta, x) = \pi(\theta, x) + u(\theta, x)$ is the total surplus, and $\lambda_{n'}$'s are the Lagrange multipliers.

Therefore, the effect of costly information acquisition amounts to a correction to the informational rent term, and its sign may vary. There are several consequences:

1. The distortion in $x(\theta)$ is different from the case when information is costless. In particular, some of the $x(\theta)$ may receive upward distortion.
2. There is no *ex ante* reason to expect that $x(\theta)$ is strictly increasing. Without the correction from costly information acquisition, monotonicity of x is guaranteed by the monotonicity of hazard rate and supermodularity of S . However, in this case, it is hard to judge whether the correction term is increasing or decreasing in θ (fixing $x(\theta)$). Therefore, the monotonicity constraint cannot be ignored in general.
3. Though the integrand of the Lagrangian equals $S(\bar{\theta}, x(\bar{\theta}))$ when $\theta = \bar{\theta}$, efficiency at the top needs not hold. When the monotonicity constraint is binding in a neighborhood of $\bar{\theta}$, distortion is created by the dependence of the Pontryagin multiplier of the monotonicity constraint on $x(\bar{\theta})$.

3.4.3 The discrete type space

In this subsection, I assume that $\bigcup_{n \in \mathcal{N}} \text{supp} E[\theta | \mathcal{F}_n]$ is a finite set, and write it as $\{\theta_1, \dots, \theta_M\}$ where $\theta_i < \theta_j$ for $i < j$. Then each information acquisition strategy $n \in \mathcal{N}$ determines a probability distribution over $\{\theta_1, \dots, \theta_M\}$. Let

$$q_i^{(n)} = \Pr(E[\theta | \mathcal{F}_n] = \theta_i), i \in \{1, \dots, M\}, n \in \mathcal{N}, \quad (3.4.5)$$

and $c(n)$ be the (expected) cost of $n \in \mathcal{N}$.

In principle, this is a special case of the general setting discussed in the previous subsection. However, when the type space is discrete, a better characterization of feasible allocations through linear programming can be obtained without resorting to Lemma 1. This characterization also sheds light on the problem of conditional optimal allocations, as the simple example in Section 5.1 will illustrate.

Consider the following problem:

1. given an allocation rule $x = (x(\theta_1), \dots, x(\theta_M))$ and an information acquisition strategy $n \in \mathcal{N}$, whether there exists a incentive compatible mechanism such that the agent finds that n is optimal and reveals his type truthfully;
2. if such a mechanism exists, what is the minimum cost (for the principal) of implementing the allocation rule and the information acquisition strategy.

This amounts to the solution of the following linear program (*LP*):

$$\begin{aligned} \min_t \quad & \sum_{i=1}^M q_i^{(n)} t_i; \\ \text{s.t.} \quad & u(\theta_i, x_i) + t_i \geq 0, \text{ for all } i; \\ & u(\theta_i, x_i) + t_i \geq u(\theta_i, x_j) + t_j, \text{ for all } i, j; \\ & \sum_i [u(\theta_i, x_i) + t_i] q_i^{(n)} - c(n) \geq \sum_i [u(\theta_i, x_i) + t_i] q_i^{(n')} - c(n'), \text{ for all } n' \in \mathcal{N}. \end{aligned}$$

In particular, an allocation rule x and information acquisition strategy n can be implemented if and only if the above linear program has a solution. Note that the IC and IR constraints are independent of the distribution over $\{\theta_1, \dots, \theta_M\}$. The following two lemmas hold without assuming the linear form $u(\theta_i, x_i) = \theta_i h(x_i)$, and therefore are interesting for their own sake.

Lemma 13. If u is increasing in θ , has increasing differences in θ and x , and each $n \in \mathcal{N}$ determines a probability distribution over $\Theta = \{\theta_1, \dots, \theta_M\}$ at cost $c(n)$, then (x, n) can be implemented if and only if there exists $s_i \in [0, \xi_i]$ for $i = 2, \dots, M$ such that

$$\sum_{i=2}^M (Q_j^{(n)} - Q_j^{(n')}) s_i \geq \sum_{i=2}^M (Q_i^{(n')} - Q_i^{(n)}) [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + c(n) - c(n'), \text{ for all } n' \in \mathcal{N}, \quad (3.4.6)$$

where

$$Q_i^{(n')} = \sum_{j=i}^M q_j^{(n')} = \Pr(\theta \geq \theta_i | n'),$$

$$\xi_i = u(\theta_i, x_i) + u(\theta_{i-1}, x_{i-1}) - u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_i).$$

Proof. The single-crossing property of u implies that the non-local IC constraints are all redundant. That u is increasing in θ , together with the IC constraints, implies that the participation constraints are redundant except the constraint for the lowest type. Therefore, the IR and IC constraints can be written as

$$u(\theta_1, x_i) + t_1 \geq 0;$$

$$u(\theta_i, x_{i-1}) - u(\theta_i, x_i) \leq t_i - t_{i-1} \leq u(\theta_{i-1}, x_{i-1}) - u(\theta_{i-1}, x_i), \text{ for } i = 2, \dots, M.$$

Clearly, the participation constraint for θ_1 will be binding at optimality. Let $s_i = t_i - t_{i-1} - [u(\theta_i, x_{i-1}) - u(\theta_i, x_i)]$ for $i = 2, \dots, M$. Then the IC constraint can be written as $0 \leq s_i \leq \xi_i$ for $i = 2, \dots, M$. Also,

$$u(\theta_i, x_i) + t_i = u(\theta_i, x_i) - u(\theta_1, x_1) + \sum_{j=2}^i [s_j + u(\theta_j, x_{j-1}) - u(\theta_j, x_j)] = \sum_{j=2}^i [s_j + u(\theta_j, x_{j-1}) - u(\theta_{j-1}, x_{j-1})].$$

Therefore, the maximum payoff associated with a particular $n' \in \mathcal{N}$ is

$$\sum_i q_i^{(n')} \sum_{j=2}^i [s_j + u(\theta_j, x_{j-1}) - u(\theta_{j-1}, x_{j-1})] - c(n') = \sum_{j=2}^M [s_j + u(\theta_j, x_{j-1}) - u(\theta_{j-1}, x_{j-1})] Q_j^{(n')} - c(n').$$

Therefore, t is feasible if and only if $s_j \in [0, \xi_j]$ and

$$\sum_{i=2}^M [s_i + u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] (Q_j^{(n)} - Q_j^{(n')}) \geq c(n') - c(n), \text{ for all } n' \in \mathcal{N}.$$

□

Not surprisingly, Lemma 2 is a special case of Lemma 1 when $u(\theta, x) = \theta h_1(x) + h_0(x)$. Specifically, Eqs. (3.4.3) and (3.4.6) are equivalent, as they both express the idea that $V(\theta_n) - c(n) \geq V(\theta_{n'}) - c(n')$ for all $n' \in \mathcal{N}$. The constraint that $s_i \in [0, \xi_i]$ is equivalent

to the monotonicity $x(\theta)$ for the following reason. The envelop theorem implies that

$$s_i = \int_{\theta_{i-1}}^{\theta_i} h_1(x(\theta))d\theta - (\theta_i - \theta_{i-1})h(x(\theta_{i-1})). \quad (3.4.7)$$

If x is monotonic, then $h(x(\theta)) \in [h(x(\theta_{i-1})), h(x(\theta_i))]$ for all $\theta \in [\theta_{i-1}, \theta_i]$, so $s_i \in [0, \xi_i]$. Conversely, given an $s_i \in [0, \xi_i]$, the fact that $\xi_i \geq 0$ implies that $x(\theta_i) \geq x(\theta_{i-1})$, and one can choose $\theta \in [\theta_{i-1}, \theta_i]$ such that

$$s_i = (\theta - \theta_{i-1})h(x(\theta_{i-1})) + (\theta_i - \theta)h(x(\theta_i)) - (\theta_i - \theta_{i-1})h(x(\theta_{i-1})).$$

Setting $x(\tilde{\theta}) = x(\theta_{i-1})$ for $\tilde{\theta} \in [\theta_{i-1}, \theta]$ and $x(\tilde{\theta}) = x(\theta_i)$ for $\tilde{\theta} \in [\theta, \theta_i]$. Then x is non-decreasing in θ , and Eq. (3.4.7) holds. Therefore, Lemma 2 and Lemma 1 are equivalent when the type space is discrete. However, the linear formulation in Lemma 2 is more tractable analytically and easier computationally, so I will use this linear programming formulation in what follows when the type space is discrete.

Lemma 14. If Eq. (3.4.6) holds in Lemma 2 for some s with $s_i \in [0, \xi_i]$, the minimum total transfer is given by

$$\sum_{i=2}^M Q_i^{(n)}[u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] - \sum_{i=1}^M q_i^{(n)}u(\theta_i, x_i) + \sum_{i=2}^M Q_i^{(n)}s_i^*. \quad (3.4.8)$$

where (s_i^*) is the vector s that minimizes $\sum_i Q_i^{(n)}s_i$ among those that satisfy Eq. (3.4.6) and lie between 0 and ξ_i .

Proof. This is obvious from the proof of Lemma 2. □

The first term in Eq. (3.4.8) is the usual "informational rent" term in the classical screening problem, the second term represents the expected utility that the agent derives from the allocation, and the third term is the additional rent that arises from the information acquisition problem. By construction, the third term is nonnegative, and its dependence on x creates additional distortion in the problem of choosing the optimal x .

As mentioned above, Lemmas 2 and 3 do not impose the linear structure on $u(\theta, x)$, nor do they assume that the posterior expectations of θ_i form a martingale (as they must do if they come from Bayesian updating), so they are applicable in more general situations.

For example, the agent can exert effort $n \in \mathcal{N}$ to improve the distribution of signal θ (in the sense of first-order-stochastic-dominance, for example.) For example, the model can describe job assignment with unobservable human capital: the agent chooses an effort $n \in \mathcal{N}$ to acquire human capital, and the outcome is a level of human capital θ . The principal offers a contract consisting pairs of job and wage. Lemmas 2 and 3 characterize the optimal contract that induces the appropriate effort of acquiring human capital and lets the agent self-select the optimal job.

3.5 Applications

In this section, I discuss two examples of the linear model developed in the previous section: binary information acquisition, and a success-failure experiment. The goal is to illustrate some qualitative effect of the costly information acquisition.

3.5.1 Binary information acquisition

Suppose that the agent can either exert effort and observe a signal θ with support $\{\theta_1, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_M\}$ with distribution $\Pr(\theta = \theta_j) = q_j$, $j \neq k$. or exert no effort and observe no signal, in which case the expectation of θ is $\theta_k \in (\theta_{k-1}, \theta_{k+1})$. To simplify notation, set $q_k = 0$. The cost of acquiring information is c . An informative mechanism is a mechanism under which the agent chooses to exert effort in equilibrium. Then in an informative mechanism, there is only one constraint associated with information acquisition in the linear program (LP):

$$\sum_{i=2}^M (Q_i - 1_{i \leq k}) s_i \geq \sum_{i=2}^M (1_{i \leq k} - Q_i) [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + c, \quad (3.5.1)$$

where $Q_i = \sum_{j=2}^M q_j$ as usual. Clearly, the allocation x can be implemented in an informative mechanism if and only if the above inequality holds when $s_i = \xi_i$ for all $i > k$ and $s_i = 0$ for all $i = k$. In other words, the following inequality holds:

$$\sum_{i=k+1}^M Q_i \xi_i \geq \sum_{i=2}^M (1_{i \leq k} - Q_i) [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + c. \quad (3.5.2)$$

Suppose the allocation rule in the optimal informative mechanism is $x = (x_1, \dots, x_M)$. In what follows, distortion is defined as the difference between x and the first best allocation given by $x_i^* \in \operatorname{argmax}_x \pi(\theta_i, x) + u(\theta_i, x)$. Assume that u and π are differentiable in x and $\pi + u$ is concave in x . Let $\Pi(c)$ be the principal's payoff in the optimal informative mechanism. Then there are four cases:

Case 1 the right hand side of Eq. (3.5.2) is negative. In this case, (LP) is solved by $s_i = 0$ for all i , and the allocation solves the classical screening problem; i.e. the second-best allocation and transfer rule in the classical screening problem are incentive compatible and optimal. In particular, there is downward distortion for all $i < M$ and no distortion for $i = M$. The principal's payoff does not depend on c in this region.

Case 2 the right hand side of Eq. (3.5.2) is zero. In this case, (LP) is solved by $s_i = 0$ for all i , but a perturbation in x_i may cause the right hand side of Eq. (3.5.2) to become positive. In this case, $\Pi'(c) < 0$, as the allocation has to be adjusted when c increases to make the right hand side of Eq. (3.5.2) remain zero. However, in this case $\Pi'(c) \geq -1$ as the principal can also choose to keep the allocation the same and increase transfers s_i to types $i > k$, and this is a one-for-one transfer of the two parties' utility.

Case 3 the right hand side of Eq. (3.5.2) is positive, but Eq. (3.5.2) holds with strict inequality. In this case the implementability of x is not a concern, so Eq. (3.5.2) is not binding in the problem of choosing x . However, Eq. (3.5.1) holds with equality with $s_i = 0$ for $i \leq k$, and thus the value function of the linear program (LP) is

$$\sum_{i=2}^k [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] - \sum_{i=1}^M q_i u(\theta_i, x_i) + c.$$

Therefore, x_{k+1}, \dots, x_M are not distorted from the first best, but x_1, \dots, x_{k-1} are downwardly distorted more compared to Case 1. Since Eq. (3.5.2) is not binding in this case, the allocation x_1, \dots, x_M does not depend on c in this region, but the principal "reimburse" the agent's cost of information through the additional transfers s_{k+1}, \dots, s_M , so $\Pi'(c) = -1$.

Case 4 Eq. (3.5.2) holds with equality. In this case, the value function of the linear program (LP) is the same as in Case 3, but the constraint Eq. (3.5.2) is binding.

After substituting in the definition of ξ_i , Eq. (3.5.2) can be rewritten as

$$-\sum_{i=2}^k (1 - Q_i)[u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + \sum_{i=k+1}^M Q_i[u(\theta_i, x_i) - u(\theta_{i-1}, x_i)] \geq c.$$

By the single-crossing property of u , the left hand side is increasing in x_i for $i \geq k+1$ and decreasing in x_i for $i \leq k-1$. Therefore, in this case there is upward distortion for x_{k+1}, \dots, x_M , and the downward distortion for x_1, \dots, x_{k-1} is bigger than in Case 2. The envelop theorem implies that $\Pi'(c) < -1$ in this case, as not only does the principal reimburse the agent's cost of information, she also distorts the allocation further to induce information acquisition.

Let x^{SB} be the second best allocation. Then the above discussion can be summarized in the following proposition:

Proposition 24. The allocation in the optimal informative mechanism is x^{SB} if

$$c \leq \sum_{i=2}^M (Q_i - 1_{i \leq k})[u(\theta_i, x_{i-1}^{SB}) - u(\theta_{i-1}, x_{i-1}^{SB})], \quad (3.5.3)$$

where it is understood that $x_k^{SB} = x_{k-1}^{SB}$. The principal's payoff does not depend on c when Eq. (3.5.3) holds. Let x^{TB} be the solution to the following program:

$$\begin{aligned} \max_x \quad & \sum_{i=1}^M q_i S(\theta_i, x_i) - \sum_{i=2}^k [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] - c & (3.5.4) \\ \text{s.t.} \quad & \sum_{i=2}^k (Q_i - 1)[u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + \sum_{i=k+1}^M Q_i [u(\theta_i, x_i) - u(\theta_{i-1}, x_i)] \geq c & (3.5.5) \end{aligned}$$

Then x^{TB} is the allocation in the optimal informative mechanism if

$$c > \sum_{i=2}^M (Q_i - 1_{i \leq k})[u(\theta_i, x_{i-1}^{TB}) - u(\theta_{i-1}, x_{i-1}^{TB})].$$

Furthermore, $x_i^{TB} \geq x_i^{FB}$ for $i = k+1, \dots, M$, with equality when Eq. (3.5.5) is not binding; $x_i^{TB} \leq x_i^{SB}$ for $i = 1, \dots, k-1$, with strict inequality if I is twice continuously differentiable and $I'(0) = 0$.

Proof. This is straight forward from the discussion of the previous section. \square

Therefore, the allocation for types $\theta_1, \dots, \theta_{k-1}$ is always downwardly distorted compared to the first best allocation. The allocation for types $\theta_{k+1}, \dots, \theta_M$, on the other hand, is weakly bigger than the second best (when information is costless), and may have no distortion or upward distortion compared to the first best. In particular, efficiency at the top may fail, though the allocation for the top type is never *downwardly* distorted. Note that the right hand side of Eq. (3.5.3) equals zero when $M = 3$, so in the two-type case $\Pi'(c) \leq -1$ always. In this sense, the two-type case is not representative.

The economic intuition is as follows. The informational rent is increasing in the allocation to the lower types. When there is cost of acquiring information, the principal has incentive to reduce the rent for θ_k and to increase the rent for types higher than the prior type θ_k , as informational rent for higher types can be captured only if the agent acquires information. To achieve this goal, the principal increases the allocation to types higher than θ_k and reduces the allocation to types lower than θ_k from the second best solution. (Note that though reducing allocation to lower types also reduces the informational rent for types $\theta_2, \dots, \theta_{k-1}$, this effect is dominated by the reduction of rent for the prior type.)

Recall the example of investing in a risky project. Proposition 4 implies that the principal invests below the first-best level when the signal is unfavorable, meaning that θ_i is below the prior. The principal invests above the second best level when the signal is favorable and c is moderate, and may invest at or above the first best level when c is sufficiently big.

One interesting observation from Proposition 4 is that the solution x^{TB} to the program Eqs. (3.5.4) and (3.5.5) does not necessarily deliver lower total surplus than x^{SB} . Compared with x^{SB} , the downward distortion is aggravated for types below the prior but mitigated or even reverted for types above the prior under x^{TB} . Therefore, it should not be surprising that the total surplus delivered by x^{TB} can exceed what x^{SB} delivers. This is indeed the case at least in a very special situation, in which there is no way to downwardly distort the allocation for low types. The next proposition states this fact formally for the general linear model with binary information acquisition:

Proposition 25. Consider the linear model with $\mathcal{N} = \{0, 1\}$. Suppose that $X \subset \mathbb{R}^+$ and $x_i^{FB} = 0$ for $i = 1, \dots, k$, and the cost of information acquisition satisfies the following

condition:

$$\sum_{i=k+1}^M Q_i [u(\theta_i, x_i^{FB}) - u(\theta_{i-1}, x_{i-1}^{FB})] < c < \sum_{i=k+1}^M Q_i [u(\theta_i, x_i^{FB}) - u(\theta_{i-1}, x_i^{FB})], \quad (3.5.6)$$

then the optimal informative mechanism implements the first best allocation.

Proof. See the appendix. □

It should be noted that the principal still prefers the second best allocation and transfer to the outcome of the optimal informative mechanism even in this case, but the second best allocation and transfer rule cannot be implemented when information acquisition is not contractible.

Now I demonstrate that the parameter set described in Proposition 5 can be nonempty in the investment example given in Section 3.2 (with a more general type space Θ). Assume that $\Theta = \{\theta_L, \theta_M, \theta_H\}$, and the agent's signal has distribution $\Pr(\theta = \theta_i) = q_i$, $i = L, M, H$. In addition,

$$I(x) = \begin{cases} x, & \text{if } x \leq x_0; \\ (1+r)x - rx_0, & \text{if } x \in (x_0, x_1]; \\ \infty, & \text{if } x > x_1, \end{cases}$$

where $r > 0$ is a given constant. The intuition is that the organization has cash x_0 with no good alternative use, and can obtain fund at shadow cost r when $x \leq x_1$, but it would be prohibitively costly to invest more than x_1 . Assume that $E[\theta]R < 1$, $\theta_H R > 1 + r$ and

$$1 < \theta_M R < \min \left\{ 1 + \frac{q_H}{q_M} \alpha (\theta_H - \theta_M) R, 1 + r \right\},$$

so that $x_M^{FB} = x_0$, $x_H^{FB} = x_1$, and $x_M^{SB} = x_L^{FB} = 0$. Eq. (3.5.6) now reads

$$q_H(\theta_H - \theta_M)\alpha x_0 R < c < q_H(\theta_H - \theta_M)\alpha x_1 R. \quad (3.5.7)$$

If this condition is satisfied, then the optimal informative mechanism implements the first best allocation. Moreover, when

$$c < q_M(\theta_M R - 1)x_0 + q_H[(\theta_H R - 1 - r)x_1 + rx_0],$$

it is optimal to induce information acquisition. (The right hand side is the total surplus generated by the information, but it is easy to check that in the optimal informative mechanism the agent gets zero informational rent, net of the cost of acquiring information.) All of the above conditions hold when, for example, when $R = 2, r = \frac{1}{2}, \alpha = \frac{1}{3}, \theta_H = 1, \theta_M = \frac{2}{3}, \theta_L = 0, q_H = \frac{1}{2}, q_M = q_L = \frac{1}{4}$, and $c \in (\frac{1}{9}x_0, \frac{1}{9}x_1)$.

One useful observation from the constraint Eq. (3.5.7) is that in order to implement the first best in the optimal informative mechanism, α cannot be too small. (When α is too small, the principal overinvests under favorable signals.) This is in contrast with the comparative statics obtained under the second best, which implies that the organization is more efficient when α is smaller. Intuitively, it is good to give the agent stake in the total surplus to motivate him to acquire information without distorting the investment.

The comparison between the optimal informative mechanism and the second best raises the following question: if there is a monitoring technology or institutional innovation that makes the effort of information acquisition contractible, is it always beneficial to contract on that? For example, the agent's cost of effort may come from hiring a consultant, but that cost may be verifiable, though the principal still lacks the expertise to interpret the consultant's report or it is prohibitively costly for her to do so. Obviously, this discussion only makes sense when n is deterministic.

When n becomes contractible, the principal can refuse to contract with the agent in Step 2 of the timing. Therefore, if the agent does not make the desired level of effort to acquire information, instead of imitating some other type and getting positive payoff, he can only get his outside option now. Therefore, Eq. (3.4.3) becomes

$$\int F^{(n)} h_1(x(\theta)) d\theta + c(n) \leq c(n'), \text{ for all } n' \neq n.$$

As this constraint is looser than Eq. (3.4.3), the principal's payoff is weakly higher in this case. Therefore, if the monitoring technology that makes n contractible is free, the principal always weakly prefers contracting on n .

On the other hand, whether contracting on n improves total surplus of the organization is a more subtle question. Part of the advantage of contracting on n for the principal is that by worsening the agent's payoff from deviation she bears less burden of the cost $c(n)$ compared to the case when n is not contractible. However, from the organization's

perspective, the full cost of information acquisition has to be deducted from the organization's total surplus. Therefore, for the organization, it boils down to the comparison of the total surplus from the allocation. Even though it is easier to implement the second best allocation when n is contractible, it is not *ex ante* clear whether the total surplus becomes higher, as the second best allocation itself suffers from downward distortion for all but the highest type. The next proposition states that in some circumstances, being able to contract on n is detrimental for the organization.

Proposition 26. Consider the discrete linear model with binary information acquisition $\mathcal{N} = \{0, 1\}$ and prior type θ_k , and let x^{SB} be the second best allocation. Assume that X is an interval, $S(\theta, x) = \pi(\theta, x) + u(\theta, x)$ is continuously differentiable and concave in x and $h_1(x)$ is convex in x . Moreover, $x_1^{SB} < x_2^{SB} < \dots < x_M^{SB}$. Then contracting on n reduces total surplus for some value of $c(1)$ if the following condition is satisfied:

$$\sum_{i \neq k} \frac{q_i S_x(\theta_i, x_i^{SB})(Q_{i+1} - 1_{i \leq k-1})(\theta_{i+1} - \theta_i)h'(x_i^{SB})}{q_i S_{xx}(\theta_i, x_i^{SB}) - Q_{i+1}(\theta_{i+1} - \theta_i)h''(x_i^{SB})} < 0.$$

Proof. See the appendix. □

The intuition of the condition is as follows. As mentioned before, when the information acquisition constraint is binding, the downward distortion for types with θ_i lower than θ_k is aggravated and the distortion for types with θ_i higher than θ_k is mitigated. The condition precisely states that the second effect dominates the first for the total surplus.

As an example, consider the problem in the previous subsection with $I(x) = x^2/2$. Then the condition in the proposition reduces to

$$\sum_{i \neq k} (\theta_{i+1} - \theta_i)^2 q_i^{-1} Q_{i+1} (Q_{i+1} - 1_{i \leq k-1}) > 0,$$

where we have substituted in the second best allocation $x_i = \theta_i R - q_i^{-1} Q_{i+1} (\theta_{i+1} - \theta_i)$, where $\theta_k - \theta_{k-1}$ should be understood as $\theta_{k+1} - \theta_{k-1}$ as θ_k represents the prior type. This condition can be satisfied when $(\theta_{i+1} - \theta_i)$ is sufficiently big for $i > k$. Therefore, contracting on n , though beneficial for the principal, may be detrimental for the organization in some circumstances.

The final proposition in this subsection characterizes the allocation in the optimal informative mechanism for a more general type space $\Theta = [\underline{\theta}, \bar{\theta}]$, showing that the economic

insights developed above extend to the more general case. Its proof relies heavily on Lagrange (geometric) multipliers, so the following assumptions are made for the strong duality theorem (existence of the multipliers) in Bertsekas (1999) to apply:

Assumption 9. $h_1(X)$ is a convex set, and the function $S(\theta, h^{-1}(\cdot))$ is concave on $h_1(X)$ for every θ .

Assumption 10. For all relevant c , there is a non-decreasing allocation $\bar{x} : \Theta \rightarrow X$ such that

$$\int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)] h_1(\bar{x}(\theta)) d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(\bar{x}(\theta)) d\theta > c.$$

Proposition 27. Assume that the total surplus $S(\theta, x) = \pi(\theta, x) + u(\theta, x)$ has increasing differences in θ and x . Let x^{FB} and x^{SB} be the first-best and second-best allocations, respectively. Let $x(\theta; c)$ be the allocation and $\Pi(c)$ be the principal's payoff in the optimal informative mechanism when the cost of information is c . Then

- i) $x(\theta-; c) \leq x^{SB}(\theta)$ for all $\theta \leq \theta_0$, and $x(\theta+; c) \geq x^{SB}(\theta)$ for $\theta \geq \theta_0$, where $x(\theta-)$ and $x(\theta+)$ denote the left and right limits at θ , respectively;
- ii) there exists $\bar{c} \geq 0$ such that $x(\theta; c) = x^{SB}(\theta)$ for all θ for $c \leq \bar{c}$;
- iii) $\Pi(c)$ is a non-increasing concave function in c ;
- iv) when $\Pi'(c) < -1$, $x(\theta+; c) \geq x^{FB}(\theta)$ for $\theta \geq \theta_0$.

Proof. See the appendix. □

3.5.2 A success-failure experiment

In this subsection, I consider such an information acquisition technology described in Example 2 of Section 2. As before, I assume that the agent's utility is $u(\theta, x) = \theta h_1(x) + h_0(x)$ for some nonnegative and strictly increasing function h_1 . Let $S(\theta, x) = \pi(\theta, x) + u(\theta, x)$ be the total surplus, which is assumed to have increasing differences in θ and x . The goal here is to compare the investment n in information acquisition in the optimal mechanism and under first best.

Given an allocation rule x and a transfer rule t , let $V(\theta) = \max_{\tilde{\theta}} u(\theta, x(\tilde{\theta})) + t(\tilde{\theta})$. The envelop theorem implies that the agent's IR and IC constraints (except for the conditions

for information acquisition) are equivalent to the following conditions:

$$V(\theta) = \int_{\underline{\theta}}^{\theta} h_1(x(\tilde{\theta}))d\tilde{\theta}; \quad (3.5.8)$$

$$x'(\theta) \geq 0. \quad (3.5.9)$$

Therefore, if the agent chooses effort $n \in [0, 1]$, his expected payoff is

$$n \int V(\theta)dF(\theta) + (1-n)V(\theta_0) - c(n) = n \int_{\underline{\theta}}^{\bar{\theta}} [1-F(\theta)]h_1(x(\theta))d\theta + (1-n) \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta))d\theta - c(n).$$

Since this function is concave in n , $n \in [0, 1]$ is optimal if and only if the following first-order condition is satisfied:

$$\int_{\underline{\theta}}^{\bar{\theta}} [1-F(\theta)]h_1(x(\theta))d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta))d\theta - c'(n) = 0. \quad (3.5.10)$$

The principal's payoff is

$$E[\pi(\theta, x(\theta)) - t(\theta)] = E[S(\theta, x(\theta)) - V(\theta)] = E[S(\theta, x(\theta))] - n \int V(\theta)dF(\theta) - (1-n)V(\theta_0).$$

Using Eqs. (3.5.8) and (3.5.10), the above function becomes

$$n \int S(\theta, x(\theta))dF(\theta) + (1-n)S(\theta_0, x(\theta_0)) - \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta))d\theta - nc'(n).$$

Let $\Pi(n)$ be the principal's expected payoff from the optimal mechanism conditional on $n \in [0, 1]$. Then

$$\Pi(n) = \max_{x(\cdot)} n \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x(\theta))dF(\theta) + (1-n)S(\theta_0, x(\theta_0)) - \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta))d\theta - nc'(n), \quad (3.5.11)$$

subject to Eqs. (3.5.9) and (3.5.10). To apply the strong duality theorem, assume that Assumptions 1 and 2 hold, with c in Assumption 2 replaced with $c'(n)$. Let n^* be the maximum of Π and x^* be an optimal allocation conditional on n^* . Then the envelop theorem implies that n^* satisfies the following first-order condition

$$\int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x^*(\theta))dF(\theta) - S(\theta_0, x^*(\theta_0)) - c'(n^*) - (n^* + \lambda)c''(c^*) = 0, \quad (3.5.12)$$

where λ is the Lagrange multiplier for the constraint Eq. (3.5.10). On the other hand, the first-best allocation $x^{FB}(\theta)$ maximizes $S(\theta, x)$ for every $\theta \in \Theta$, and the first best

information acquisition satisfies

$$n^{FB} \in \operatorname{argmax}_n n \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x^{FB}(\theta)) dF(\theta) + (1-n)S(\theta_0, x^{FB}(\theta_0)) - c(n).$$

Hence, n^{FB} satisfies the following first-order condition:

$$\int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x^{FB}(\theta)) dF(\theta) - S(\theta_0, x^{FB}(\theta_0)) - c'(n^{FB}) = 0. \quad (3.5.13)$$

The following lemma gives a sufficient condition for n^* to be less than or equal to n^{FB} and some characterization of the allocation in the optimal mechanism.

Lemma 15. If $\lambda \geq 0$, then i) $x^*(\theta_0) = x^{FB}(\theta_0)$; ii) $x^*(\theta-) \leq x^{FB}(\theta)$ for all $\theta \leq \theta_0$ and $x^*(\theta+) \geq x^{FB}(\theta)$ for all $\theta \geq \theta_0$, where $x^*(\theta-)$ and $x^*(\theta+)$ are the left and right limits of x^* at θ , respectively; and iii) $n^* \leq n^{FB}$, with strict inequality if c is strictly convex.

Proof. See the appendix. □

To understand the intuition of the lemma, recall that when $\lambda > 0$ the constraint Eq. (3.5.10) is equivalent to the same expression with the equality replaced with "≥". In this case, the agent would acquire less information than the desired level n^* without incentives, so the principal has to concede part of the benefit of information to the agent. This implies that the marginal benefit of information for the principal is less than its first-best level. Therefore, less information is acquired in the optimal mechanism. The second assertion in the lemma roughly says that there is a downward distortion for $\theta < \theta_0$ and upward distortion for $\theta > \theta_0$, and this statement is precise when F is absolutely continuous with respect to the Lebesgue measure.

The problem with Lemma 4 is that it is stated in terms of the Lagrange multiplier λ , whose sign is hard to determine without solving the whole program Eqs. (3.5.11), (3.5.9) and (3.5.10). The next proposition gives a sufficient condition which can be checked without solving the program.

Proposition 28. Assume that the p.d.f. of F is bounded away from zero. Then $n^* \leq n^{FB}$ if the following condition is satisfied:

$$\int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)] h_1(x^{FB}(\theta)) d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x^{FB}(\theta)) d\theta < c'(n^{FB}).$$

Moreover, under the above condition, $n^* < n^{FB}$ if c is strictly convex.

Proof. See the appendix. □

The condition given in this proposition has a very intuitive interpretation: insufficient information (compared to first-best) is produced in the optimal mechanism if the agent prefers acquiring less information than n^{FB} under the allocation rule x^{FB} . What is not clear from Proposition 8 (and Lemma 4) is that whether $n^* \leq n^{FB}$ when the conditions stated in the proposition or the lemma do not hold. In that case, $x^*(\theta_0)$ may be distorted, which makes the comparative statics more complicated. This is left for future research.

3.6 Conclusion

The main contribution of this paper is to set up a principal-agent model that allows for general endogenous information acquisition, and to analyze the cases of binary information acquisition and success-failure experiment. It is shown that results obtained in previous works such as Crémer *et al* (1998b) and Lewis and Sappington (1997) arise in the general model too, but more interesting phenomena can also happen. The consideration of the success-failure experiment in this context and the sufficient condition for less-than-efficient information acquisition is an innovation of this paper.

As the set up of the model is relatively general, there can potentially be a wide range of applications. These include investing in risky projects, human capital acquisition, regulation and procurement among others. In particular, when the information acquisition is binary or of the form of a success-failure experiment, the analysis in Section 5 delivers interesting comparative statics, that can be tested empirically or compared to other models.

From a theoretical perspective, the analysis here can be extended in several directions. First, it would be interesting to have an example with contingent information acquisition (optimal stopping) which admits a closed-form solution. This will provide new economic insights in a broader range of applications, as well as check the robustness of the predictions in Section 5.

A second extension is to the case with more than one agents. Though optimal auctions have been studied by previous works such as Shi (2008), more interesting multi-agent mechanisms with endogenous information acquisition are yet to be explored. For example, the agent may engage in a tournament or compete for the scarce resources, and it would be interesting to see how the agents' effort of acquiring information and the distortion in the allocation depend on the number of agents.

In many applications, θ is interpreted as the probability that some event happens. For example, in the investment problem considered in Section 3.2, θ is the probability that the risky project succeeds. However, in reality there might be multiple outcomes and decisions must be made. For example, the principal needs to allocate resources in more than one risky projects. In this case, θ becomes an element of a simplex, but the agent's (and the principal's) payoff is still linear in θ . This calls for a generalized framework which allow both θ and x to be multi-dimensional. The question for future research is whether there exists some appropriate single-crossing condition to guarantee the convexity of the agent's value, and whether the comparative statics analogous to those in Section 5 can be obtained when information acquisition technology is as in that section.

Finally, it is worth considering the repeated interaction between the principal and the agent. Besides the standard constraints that relate the implementable outcomes to the patience of the parties in the relational contract literature, the repeated game can also give the agent additional incentives to acquire information. For simplicity, consider the case when the state of nature θ evolves as a Markov process, and the principal's continuation strategy can depend on the state of nature that agent reported in the past but not on the actual past state of nature. Then the ergodic distribution of the agent's report is determined by the transition matrix of the Markov process, assuming that the agent always chooses the information acquisition action that the principal recommends and tells the truth. Therefore, any deviation in the empirical distribution of the agent's reports from the ergodic distribution can be punished by the principal, and this makes it more difficult for the agent to choose another level of information and imitate some type. For example, in the binary information acquisition case, if the principal recommends the agent to acquire information, the agent can always choose not to do so and imitate some type θ' that leads to the highest expected payoff. In the repeated interaction, however, the principal can detect this because the probability that the state of nature is always θ' is

very low. However, studying the repeated interaction carefully requires more machinery and is beyond the scope of this paper.

3.7 Appendix

Proof of Proposition 5. Let x be the allocation in the optimal informative mechanism. Suppose the right hand side of Eq. (3.5.2) is not positive at optimality. Then $x_i \leq x_i^{FB}$ for all i . In particular, $x_i = 0$ for all $i \leq k$. The single-crossing property of u implies that $u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1}) \leq u(\theta_i, x_{i-1}^{FB}) - u(\theta_{i-1}, x_{i-1}^{FB})$ for $i = 2, \dots, M$. Therefore,

$$\sum_{i=2}^M (1_{i \leq k} - Q_i) [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + c \geq - \sum_{i=k+1}^M Q_i [u(\theta_i, x_{i-1}^{FB}) - u(\theta_{i-1}, x_{i-1}^{FB})] + c > 0,$$

a contradiction. Therefore, the right hand side of Eq. (3.5.2) is positive, and x solves the following program:

$$\begin{aligned} \max_x \quad & \sum_i q_i S(\theta_i, x_i) - \sum_{i=2}^k [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] - c \\ \text{s.t.} \quad & \sum_{i=2}^k (Q_i - 1) [u(\theta_i, x_{i-1}) - u(\theta_{i-1}, x_{i-1})] + \sum_{i=k+1}^M Q_i [u(\theta_i, x_i) - u(\theta_{i-1}, x_i)] \geq c. \end{aligned}$$

Clearly, $x = x^{FB}$ maximizes the objective function, ignoring the constraint. Eq. (3.5.6) implies that the constraint is satisfied by x^{FB} .

Proof of Proposition 6. Under the assumption, the second best allocation satisfies the following first order condition:

$$q_i S_x(\theta_i, x_i) - Q_{i+1}(\theta_{i+1} - \theta_i) h'(x_i) = 0, \text{ for } i \neq k.$$

Let c_0 be the maximum cost of information acquisition such that the second best allocation is implementable when n is not contractible. Then when c is in a neighborhood of c_0 , the allocation in the optimal informative mechanism satisfies the following first order condition:

$$q_i S_x(\theta_i, x_i) - [(1 - \lambda)Q_{i+1} + \lambda 1_{1 \leq k-1}] (\theta_{i+1} - \theta_i) h'(x_i) = 0, \text{ for } i \neq k,$$

where λ is the Lagrange multiplier and $\lambda > 0$ when $c > c_0$. The implicit function theorem implies that

$$\frac{\partial x_i}{\partial \lambda} = - \frac{Q_{i+1} - 1_{1 \leq k-1} (\theta_{i+1} - \theta_i) h'(x_i)}{q_i S_{xx}(\theta_i, x_i) - [(1 - \lambda)Q_{i+1} + \lambda 1_{1 \leq k-1}] (\theta_{i+1} - \theta_i) h''(x_i)}.$$

This implies that for λ close to 0, $\partial c/\partial \lambda$ is positive and bounded away from both 0 and ∞ . Therefore, there exists $\delta_1 > 0$ such that the map from λ to c is a local diffeomorphism for $\lambda \in (0, \delta_1)$. In particular, for c close to c_0 , the monotonicity constraint for x_i is not binding and x_i satisfies the above first order condition. The condition in the proposition states that $\partial \sum_i q_i S(\theta_i, x_i)/\partial \lambda > 0$ when $\lambda = 0$. This implies that there exists $\delta_2 \in (0, \delta_1)$ such that the total surplus in the optimal informative mechanism is increasing in λ for $\lambda \in [0, \delta_2)$ and thus increasing in c for $c \in [c_0, c_0 + \epsilon)$ for some $\epsilon > 0$. On the other hand, by the strict monotonicity of x_i^{SB} , the rent for the prior type is positive, so for c close to c_0 , the second best allocation is still implementable when n is contractible.

Lemma 16. Let F be a probability distribution over $\Theta = [\underline{\theta}, \bar{\theta}]$, $X \subset \mathbb{R}$, h a strictly increasing function on X and $S : \Theta \times X \rightarrow \mathbb{R}$ a function with increasing differences. Fix $\theta_0 \in (\underline{\theta}, \bar{\theta})$. For $\mu_1, \mu_2 \in \mathbb{R}$, let $x(\cdot; \mu_1, \mu_2) : \Theta \rightarrow X$ be a solution to the following program:

$$\begin{aligned} \max_{x(\cdot)} \quad & \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x(\theta)) dF(\theta) + \mu_1 \int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)] h(x(\theta)) d\theta + \mu_2 \int_{\underline{\theta}}^{\theta_0} h(x(\theta)) d\theta; \\ \text{s.t.} \quad & x \text{ is weakly increasing.} \end{aligned}$$

Then for any $\mu_1, \mu_2, \mu'_1, \mu'_2 \in \mathbb{R}$,

- i) if $\mu'_2 \leq \mu_2$ and $\mu'_1 - \mu_1 \leq \mu_2 - \mu'_2$ with at least one strict inequality, then $x(\theta-; \mu'_1, \mu'_2) \leq x(\theta; \mu_1, \mu_2)$ for all $\theta \leq \theta_0$; and
- ii) if $\mu'_1 > \mu_1$, then $x(\theta+; \mu'_1, \mu'_2) \geq x(\theta; \mu_1, \mu_2)$ for all $\theta \geq \theta_0$.

Proof. Suppose $x(\theta-; \mu'_1, \mu'_2) > x(\theta; \mu_1, \mu_2)$ for some $\theta \leq \theta_0$. Let

$$\begin{aligned} \theta_1 &= \inf\{\tilde{\theta} \in [\underline{\theta}, \theta] : x(\tilde{\theta}; \mu'_1, \mu'_2) > x(\tilde{\theta}; \mu_1, \mu_2) \text{ for all } \hat{\theta} \in [\tilde{\theta}, \theta]\}; \\ \theta_2 &= \sup\{\hat{\theta} \in [\theta, \theta_0] : x(\hat{\theta}; \mu'_1, \mu'_2) > x(\hat{\theta}; \mu_1, \mu_2) \text{ for all } \hat{\theta} \in [\theta, \hat{\theta}]\}. \end{aligned}$$

Then by assumption $\theta_1 < \theta \leq \theta_2 \leq \theta_0$ and $x(\bar{\theta}; \mu'_1, \mu'_2) > x(\bar{\theta}; \mu_1, \mu_2)$ for all $\bar{\theta} \in (\theta_1, \theta_2)$.

Let

$$\underline{x} = \begin{cases} x(\theta_1+; \mu_1, \mu_2), & \text{if } x(\theta_1; \mu'_1, \mu'_2) \leq x(\theta_1+; \mu_1, \mu_2); \\ x(\theta_1; \mu_1, \mu_2), & \text{if } x(\theta_1; \mu'_1, \mu'_2) > x(\theta_1+; \mu_1, \mu_2); \end{cases}$$

$$\bar{x} = \begin{cases} x(\theta_2-; \mu'_1, \mu'_2), & \text{if } x(\theta_2; \mu_1, \mu_2) \geq x(\theta_2-; \mu'_1, \mu'_2); \\ x(\theta_2; \mu'_1, \mu'_2), & \text{if } x(\theta_2; \mu_1, \mu_2) < x(\theta_2-; \mu'_1, \mu'_2). \end{cases}$$

Consider the following program (P')

$$\begin{aligned} \max_{x: [\theta_1, \theta_2] \rightarrow X} \quad & \int_{\theta_1}^{\theta_2} S(\theta, x(\theta)) dF(\theta) + \bar{\mu}_1 \int_{\theta_1}^{\theta_2} [1 - F(\theta)] h(x(\theta)) d\theta + \bar{\mu}_2 \int_{\theta_1}^{\theta_2} h(x(\theta)) d\theta, \\ \text{s.t.} \quad & x \text{ is weakly increasing;} \\ & \underline{x} \leq x(\theta) \leq \bar{x} \text{ for all } \theta \in [\theta_1, \theta_2], \end{aligned}$$

where the first integral in the objective includes the end point θ_1 if and only if $x(\theta_1; \mu'_1, \mu'_2) > x(\theta_1+; \mu_1, \mu_2)$ and includes the end point θ_2 if and only if $x(\theta_2; \mu_1, \mu_2) < x(\theta_2-; \mu'_1, \mu'_2)$. Define $x_1 : [\theta_1, \theta_2] \rightarrow X$ as follows: $x_1(\theta) = x(\theta; \mu_1, \mu_2)$ for $\theta \in (\theta_1, \theta_2)$; $x_1(\theta_1) = x(\theta_1; \mu_1, \mu_2)$ if θ_1 is included in the first integral in the objective of (P') and $x_1(\theta_1) = x(\theta_1+; \mu_1, \mu_2)$ otherwise; $x_1(\theta_2) = x(\theta_2; \mu_1, \mu_2)$ if θ_2 is included in the first integral in the objective of (P') and $x_1(\theta_2) = x(\theta_2-; \mu_1, \mu_2)$ otherwise. Define x_2 the same way for $x(\cdot; \mu'_1, \mu'_2)$. Then by construction, x_1 solves the program (P') when $(\bar{\mu}_1, \bar{\mu}_2) = (\mu_1, \mu_2)$, and x_2 solves the program (P') when $(\bar{\mu}_1, \bar{\mu}_2) = (\mu'_1, \mu'_2)$. Therefore,

$$\begin{aligned} & \int_{\theta_1}^{\theta_2} S(\theta, x_1(\theta)) dF(\theta) + \mu_1 \int_{\theta_1}^{\theta_2} [1 - F(\theta)] h(x_1(\theta)) d\theta + \mu_2 \int_{\theta_1}^{\theta_2} h(x_1(\theta)) d\theta \geq \\ & \int_{\theta_1}^{\theta_2} S(\theta, x_2(\theta)) dF(\theta) + \mu_1 \int_{\theta_1}^{\theta_2} [1 - F(\theta)] h(x_2(\theta)) d\theta + \mu_2 \int_{\theta_1}^{\theta_2} h(x_2(\theta)) d\theta; \\ & \int_{\theta_1}^{\theta_2} S(\theta, x_2(\theta)) dF(\theta) + \mu'_1 \int_{\theta_1}^{\theta_2} [1 - F(\theta)] h(x_2(\theta)) d\theta + \mu'_2 \int_{\theta_1}^{\theta_2} h(x_2(\theta)) d\theta \geq \\ & \int_{\theta_1}^{\theta_2} S(\theta, x_1(\theta)) dF(\theta) + \mu'_1 \int_{\theta_1}^{\theta_2} [1 - F(\theta)] h(x_1(\theta)) d\theta + \mu'_2 \int_{\theta_1}^{\theta_2} h(x_1(\theta)) d\theta. \end{aligned}$$

Adding the inequalities up, one obtains that

$$\int_{\theta_1}^{\theta_2} [(\mu'_1 - \mu_1)(1 - F(\theta)) + (\mu'_2 - \mu_2)][h(x_2(\theta)) - h(x_1(\theta))] d\theta \geq 0$$

By assumption, $(\mu'_1 - \mu_1)(1 - F(\theta)) + (\mu'_2 - \mu_2) < 0$ and $h(x_2(\theta)) > h(x_1(\theta))$ for all $\theta \in (\theta_1, \theta_2)$, a contradiction. Therefore, $x(\theta-; \mu'_1, \mu'_2) \leq x(\theta; \mu_1, \mu_2)$ for all $\theta \leq \theta_0$.

ii) can be proved in the same way. \square

Proof of Proposition 7. Let $\theta_0 = E[\theta]$, the prior of θ . By Lemma 1, the allocation in the optimal mechanism solves the following program:

$$\begin{aligned} \Pi(c) = \max_{x(\cdot)} & \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x(\theta)) dF(\theta) - \int_{\underline{\theta}}^{\bar{\theta}} (1 - F(\theta)) h_1(x(\theta)) d\theta; \\ \text{s.t.} & \quad x \text{ is non-decreasing;} \\ & \quad \int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)] h_1(x(\theta)) d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta)) d\theta \geq c. \end{aligned}$$

I will refer to the second constraint as the information inducing constraint. Using the method of Lagrange multiplier, Therefore, $x(\theta; c)$ solves

$$\begin{aligned} \Pi(c) = \max_{x(\cdot)} & \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x(\theta)) dF(\theta) - (1 - \lambda) \int_{\underline{\theta}}^{\bar{\theta}} (1 - F(\theta)) h_1(x(\theta)) d\theta - \lambda \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta)) d\theta - \lambda c; \\ \text{s.t.} & \quad x \text{ is non-decreasing,} \end{aligned}$$

where $\lambda \geq 0$ is the Lagrange multiplier. The envelop theorem implies that $\Pi'(c) = -\lambda \leq 0$. Moreover, Lemma 5 implies that $x(\theta-; c)$ is decreasing in λ for $\theta \leq \theta_0$ and $x(\theta+; c)$ is increasing in λ for $\theta \geq \theta_0$. The left hand side of the information inducing constraint is increasing in $x(\theta)$ for $\theta > \theta_0$ and decreasing in $x(\theta)$ for $\theta \leq \theta_0$. Since the discontinuities of a monotonic function have Lebesgue measure zero, the left hand side of the information inducing constraint is increasing in λ . Therefore, λ is increasing in c . This implies that $\Pi'(c)$ is decreasing in c , $x(\theta-; c)$ is decreasing in c for $\theta \leq \theta_0$ and $x(\theta+; c)$ is increasing in c for $\theta \geq \theta_0$. The second best allocation coincides with the allocation in the optimal informative mechanism when $\lambda = 0$, which happens when

$$c \leq \int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)] h_1(x^{SB}(\theta)) d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x^{SB}(\theta)) d\theta.$$

Lemma 5 also implies the comparison between $x(\cdot; c)$ and x^{SB} stated in the proposition. Finally, x^{FB} solves the program in Lemma 5 with $\mu_1 = \mu_2 = 0$, so Lemma 5 implies that $x(\theta+; c) \geq x^{FB}(\theta)$ for $\theta \geq \theta_0$ when $\lambda = -\Pi'(c) > 1$.

Proof of Lemma 4. The result is trivial if $n^* = 0$. In what follows assume that $n^* > 0$. Lemma 5 implies that $x^*(\theta-) \leq x^{FB}(\theta)$ for all $\theta \leq \theta_0$ and $x^*(\theta+) \geq x^{FB}(\theta)$ for $\theta \geq \theta_0$. Therefore, $x^*(\theta_0-) \leq x^{FB}(\theta_0) \leq x^*(\theta_0+)$, which implies that $x^*(\theta_0) = x^{FB}(\theta_0)$, as $x^{FB}(\theta_0)$ maximizes $S(\theta_0, x)$.

The second claim follows directly from Lemma 5.

Finally, by Eq. (3.5.12),

$$c'(n^*) = \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x^*(\theta))dF(\theta) - S(\theta_0, x^*(\theta_0)) - (n^* + \lambda)c''(n^*)$$

Applying the definition of x^{FB} to the first term, the result that $x^{FB}(\theta_0) = x^*(\theta_0)$ to the second, and the fact that $(n^* + \lambda)c''(n^*) \geq 0$ to the third, one obtains that

$$c'(n^*) \leq \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x^{FB}(\theta))dF(\theta) - S(\theta_0, x^{FB}(\theta_0)) = c'(n^{FB}),$$

with strict inequality if c is strictly convex and thus $c''(n^*) > 0$. Therefore, the second assertion follows from the convexity of c .

Proof of Proposition 8. Suppose $n^* > n^{FB}$. Clearly, $\lambda > 0$ if and only if the left hand side of Eq. (3.5.10) is less than the right hand side, when (x^{TB}, n^*) is substituted in, for all solution x^{TB} to the following program:

$$\begin{aligned} \max_{x(\cdot)} \quad & n^* \int_{\underline{\theta}}^{\bar{\theta}} S(\theta, x(\theta))dF(\theta) + (1 - n^*)S(\theta_0, x(\theta_0)) - \int_{\underline{\theta}}^{\theta_0} h_1(x(\theta))d\theta; \\ \text{s.t.} \quad & x'(\theta) \geq 0. \end{aligned}$$

Lemma 5 can be applied to any x^{TB} to show that $x^{TB}(\theta_0) = x^{FB}(\theta_0)$, $x^{TB}(\theta-) \leq x^{FB}(\theta)$ for all $\theta < \theta_0$. Clearly $x^{TB}(\theta) = x^{FB}(\theta)$, a.s. F , and thus almost everywhere with respect to the Lebesgue measure. Therefore,

$$\int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)]h_1(x^{TB}(\theta))d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x^{TB}(\theta))d\theta \leq \int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)]h_1(x^{FB}(\theta))d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x^{FB}(\theta))d\theta.$$

Combining this with the condition in the proposition, one obtains that

$$\int_{\underline{\theta}}^{\bar{\theta}} [1 - F(\theta)]h_1(x^{TB}(\theta))d\theta - \int_{\underline{\theta}}^{\theta_0} h_1(x^{TB}(\theta))d\theta < c'(n^{FB}) < c'(n^*),$$

for all x^{TB} , which implies that $\lambda > 0$. By Lemma 4, $n^* \leq n^{FB}$, a contradiction. Therefore, $n^* \leq n^{FB}$. Note that $\lambda > 0$ even when $n^* = n^{FB}$, so $n^* < n^{FB}$ when c is strictly convex.

Chapter 4

References

- Abreu, D., D. Pearce, and E. Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041-1063.
- Aghion, P., and J. Tirole. (1997). "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105, 1-29.
- Athey, S., and K. Bagwell (2001), "Optimal Collusion with Private Information," *The RAND Journal of Economics*, 82 (3), 428-485.
- Athey, S., K. Bagwell, and C. Sanchirico (2004), "Collusion and Price Rigidity," *Review of Economic Studies*, 71(2), 317-349.
- Baily, M. N., C. Hulten, D. Campbell, T. Bresnahan, and R. E. Caves (1992), "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity. Microeconomics*, 1992, 187-267.
- Baker, G., R. Gibbons, and K. Murphy (1994), "Subjective Performance Measures in Optimal Incentive Contract," *Quarterly Journal of Economics*, 109(4), 1125-1156.
- Baker, G., R. Gibbons, and K. Murphy (2002), "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics*, 117(1), 39-84.
- Baron, D., and R. Myerson (1983), "Regulating a Monopolist with Unknown Cost," *Econometrica*, 50, 911-930.
- Barron, D. (2013), "Attaining Efficiency with Imperfect Public Monitoring and Markov Adverse Selection," mimeo.
- Bergemann, D., and J. Välimäki. (2002). "Information Acquisition and Efficient Mechanism Design." *Econometrica*, 70(3), 1007-1033.
- Bertsekas. (1999). *Nonlinear Programming*. Belmont: Athena Scientific.
- Biais, B., T. Mariotti, C. Plantin, and J-C. Rochet (2007), "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies*, 74(2), 345-390.
- Biais, B., T. Mariotti, and J-C. Rochet (2011), "Dynamic Financial Contracting,"

mimeo.

Bloom, N. (2009), "Impacts of Uncertainty Shocks," *Econometrica*, 77(3), 623-685.

Bonatti A., and J. Horner (2011), "Collaborating," *The American Economic Review*, 101(2), 632-663.

Chassang, S. (2010), "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts," *American Economic Review*, 100(1), 448-465.

Crawford, V. and J. Sobel. (1982). "Strategic Information Transmission." *Econometrica*, 50, 1431-1451.

Crémer, J., and F. Khalil. (1992). "Gathering Information Before Signing a Contract." *The American Economic Review*, 82(3), 566-578.

Crémer, J., F. Khalil and J.-C. Rochet. (1998). "Strategic Information Gathering Before a Contract is Offered." *Journal of Economic Theory*, 81, 163-200.

Crémer, J., F. Khalil and J.-C. Rochet. (1998). "Contracts and Productive Information Gathering." *Games and Economic Behavior*, 25, 174-193.

Demarzo, P.M., and Sannikov, Y. (2006), "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 16(6), 2681-2724.

Demski, J. S., and D. E. M. Sappington. (1987). "Delegated Expertise." *Journal of Accounting Research*, 25(1), 68-89.

Dudley R.M. (2002), *Real Analysis and Probability*, Cambridge University Press, 2002.

Eeckout, J. (2006), "Minorities and Endogenous Segregation," *Review of Economic Studies*, 73(1), 31-53.

Escobar, J., and J. Toikka (2013), "Efficiency in Games with Markovian Private Information," *Econometrica*, 81(5), 1887-1934.

Foster, L., J. Haltiwanger, and C. Syverson (2008), "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1), 394-425.

Futia, C.A. (1982), "Invariant Distributions and the Limiting Behavior of Markovian Economic Models," *Econometrica*, 50(2), 377-408.

Fudenberg, D., and D. Levine (1994), "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62, 103-135.

Fudenberg D., D. Levine, and E. Maskin (1994), "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5), 997-1039.

Gary-Bobo, R., and Y. Spiegel (2006), "Optimal State-Contingent Regulation under Limited Liability," *The RAND Journal of Economics*, 37(2), 431-448.

Gibbons, R. and R. Henderson (2012), "What Do Managers Do?" *Handbook of Organizational Economics*, Princeton University Press.

Green, E.J. (1987), "Lending and the Smoothing of Uninsurable Income," *Contractual Arrangements and International Trade*, Minnesota Press, Minnesota.

- Halac M. (2012), "Relational Contracts and the Value of Relationships," *The American Economic Review*, 102(2), 750-779.
- Halac M. (2013), "Investing in a Relationship," mimeo.
- Halac, M., N. Kartik, and Q. Liu (2013), "Optimal Contracts for Experimentation," mimeo.
- Hairer, M. (2010), "Convergence of Markov Processes," mimeo.
- Hastings, D. (1999), "Lincoln Electric's Harsh Lessons from International Expansion," *Harvard Business Review*, 77(8), 162-178.
- Hörner J., S. Takahashi, and N. Vieille (2013), "Truthful Equilibria in Dynamic Bayesian Games," mimeo.
- Hsieh, C-T., P. J. Klenow (2009), "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124(4), 1403-1448.
- Inderst, R. and M. Klein. (2007). "Innovation, Endogenous Overinvestment, and Incentive Pay." *The RAND Journal of Economics*, 38(4), 881-904.
- Jenter D., and F. Kanaan (2013), "CEO Turnover and Relative Performance Evaluation," forthcoming in *The Journal of Finance*
- Jovanovic, B. (1979), "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87(5), 972-990.
- Jovanovic, B. (1982), "Selection and the Evolution of Industry," *Econometrica*, 50(3), 649-670.
- Kocherlakota, N.R. (1996), "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies*, 63, 595-609.
- Kranton, R. (1996), "The Formation of Cooperative Relationships," *Journal of Law, Economics, and Organizations*, 12(1), 214-233.
- Laffont J., and J. Tirole (1986), "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94, 614-641.
- Lambert, R. A. (1986). "Executive Effort and Selection of Risky Projects." *The RAND Journal of Economics*, 17(1), 77-88.
- Levin, J. (2003), "Relational Incentive Contracts," *American Economic Review*, 93(3), 835-857.
- Lewis, T. R. and D. E. M. Sappington. (1997). "Information Management in Incentive Problems." *The Journal of Political Economy*, 105(4), 796-821.
- Macchiavello, R. and A. Morjaria (2012), "Value of Relationships: Evidence from a Supply Shock to Kenyan Flower Exporters," Mimeo.
- Macleod, B., and J. Malcomson (1989), "Implicit Contracts, Incentive Compatibility, and Involuntary Employment," *Econometrica*, 57(2), 447-480.
- McAdams, D. (2011), "Performance and Turnover in a Stochastic Partnership," *American Economic Journal: Micro*, 3(4), 107-142.

- Meyn, S., and R.L. Tweedie (2009), *Markov Chains and Stochastic Stability*, 2nd Edition, Cambridge University Press.
- Mirrlees, J. A. (1975), "The Theory of Moral Hazard and Unobservable Behaviour, Part 1," Mimeo, Oxford. (Published in 1999 in the *Review of Economic Studies*, 66, 3-21.)
- Phelan, C., and R.M. Townsend (1991), "Computing Multi-period, Information Constrained Optima," *Review of Economic Studies*, 58(5), 853-881.
- Riordan, M. H., and D. E. M. Sappington (1988), "Optimal Contracts with Public ex post Information," *Journal of Economic Theory*, 45, 189-199.
- Rogerson, W.P. (1985), "Repeated Moral Hazard," *Econometrica*, 53(1), 69-76.
- Shivdasani, A., and D. Yermack (1999), "CEO Involvement in the Selection of New Board Members: An Empirical Analysis," *The Journal of Finance*, 54(5), 1829-1853.
- Sannikov, Y. (2007), "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica*, 75(5), 1285-1329.
- Sannikov, Y. (2008), "A Continuous-Time Version of the Principal-Agent Problem," *Review of Economic Studies*, 75, 957-984.
- Schoar A., and E. Washington (2011), "Are the Seeds of Bad Performance Sown in Good Times?" mimeo.
- Shapiro C., and J. E. Stiglitz (1984), "Equilibrium Unemployment as a Worker Discipline Device," *The American Economic Review*, 74(3), 433-444.
- Shi, X. (2008). "Optimal Auction with Information Acquisition." *Working Paper*.
- Song Z., K. Storesletten, and F. Zilibotti (2011), "Growing Like China," *The American Economic Review*, 101(1), 196-233.
- Spear, S. E., and S. Srivastava (1987), "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 54(4), 599-617.
- Syverson, C. (2004), "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112(6), 1181-1222.
- Syverson, C. (2011), "What Determines Productivity?" *Journal of Economic Literature*, 49(2), 326-365.
- Thomas, J., and T. Worrall (1988), "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55, 541-554.
- Thomas, J., and T. Worrall (1990), "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-agent Problem," *Journal of Economic Theory*, 51(2), 367-390.
- Zermeno, L. (2011). "Information Acquisition, Authority and Accountability." *Working paper*.
- Zhu, J.Y. (2013), "Optimal Contracts with Shirking," *Review of Economic Studies*, 80(2), 813-839.