In Search of a Unicorn: Dynamic Agency with Endogenous Investment Opportunities*

Felix Zhiyu Feng[†] Robin Yifan Luo[‡] Beatrice Michaeli[§]

September 2023

*We thank Mark Lang (Editor) and an anonymous referee for helpful comments and suggestions. We also thank Tim Baldenius, Bruno Biais, Matthew Bloomfield, Judson Caskey, Edwige Cheynel, Henry Friedman, Robert Göx, Mirko Heinle, Daniele Macciocchi, Erik Madsen, Mark Maffett, Simon Mayer, Delphine Samuels, Richard Saouma, Haresh Sapra, Tarik Umar, and participants at the Organizational Economics Workshop, 33rd Stony Brook International Conference on Game Theory, HKUST Accounting Symposium, Financial Markets and Corporate Governance Conference, CICF, CFRIC, APAD, and RSFE.

[†]Felix Zhiyu Feng is at Foster School of Business, University of Washington and can be reached at ffeng@uw.edu.

[‡]Robin Yifan Luo is at Foster School of Business, University of Washington and can be reached at yifanluo@uw.edu.

[§]Beatrice Michaeli is at Anderson School of Management, University of California, Los Angeles and can be reached at beatrice.michaeli@anderson.ucla.edu

In Search of a Unicorn: Dynamic Agency with Endogenous Investment Opportunities

Abstract: We study the optimal dynamic contract that provides incentives for an agent (e.g., SPAC sponsor, VC general partner, CTO) to exploit investment opportunities/targets that arrive randomly over time via a costly search process. The agent is privy to the arrival as well as to the quality of the target and can take advantage of this for rent extraction during the search stage and the ensuing production stage. The optimal contract provides the agent with incentives for timely and truthful reporting via a time-varying threshold for investment and an internal charge for the time spent on search. In the equilibrium, as time elapses, the charge becomes progressively higher while the investment threshold progressively lower, resulting in overinvestment at a time-varying degree. Our model generates empirically testable predictions regarding internal innovations and project selection as well as external investments such as M&As, hedge fund activism, VC investing, and SPACs, linking the degree of overinvestment to observable firm and industry characteristics.

JEL classification: G31, D82, D83, D86, M11

Keywords: dynamic agency, endogenous search, project selection, overinvestment

1 Introduction

Investment decisions are an important element of corporate operations. A growing stream of literature (e.g., Bernardo, Cai, and Luo, 2001; Baldenius, 2003; Heinle, Ross, and Saouma, 2014; Bastian-Johnson, Pfeiffer, and Schneider, 2017) studies these decisions in settings where the quality of the considered opportunities is only observable to the managers in charge, thereby emphasizing adverse selection as a main friction. The typical advice for achieving goal congruence in such setting is to impose a static and sufficiently high hurdle of acceptable quality, which results in ex post underinvestment. Notably, a recurring theme in this literature is that the presence of a business opportunity is given, and the timing of the (one-shot) decision to invest or not is fixed.

In practice, both external and internal investment opportunities often arise and materialize over time and may demand a significant commitment of resources for their discovery and selection. For instance, Sponsors of Special Purpose Acquisition Companies (SPACs) spend months or years seeking out the most valuable private enterprises—commonly referred to as "unicorns"—with which to merge. Chief Executive Officers (CEOs) embark on an enduring quest to identify acquisition targets for maintaining corporate growth. General Partners and Associates of venture capital (VC) funds scout for nascent startups primed for investment. Managers of activist hedge funds engage in a pursuit of undervalued firms for intervention. Chief Technology Officers (CTOs) from the high-tech sector invest their time and company's resources into researching and developing pioneering technologies. Chief Operating Officers (COOs) and Divisional Managers consider multiple projects—a process that unfolds over an extensive span of time—before singling out the one with the most auspicious prospects.

Furthermore, the decisions of when to invest in a given opportunity or end the search process are neither exogenous nor made in a static environment: managers and investors are aware that future opportunities may arrive if they do not invest immediately, but do not necessarily know ex ante how many opportunities are available and what their quality will be. Meanwhile, there is no need to abide by a fixed schedule that dictates the time at which an investment decision has to be made. Consequently, the criteria for investment as well as the decision whether to continue or abandon the search can all be dynamic and non-stationary, depending on the history of the managers' search.

In this paper we study the implications of endogenous and dynamic search for opportunities on investment decisions, and demonstrate that these features lead to novel predictions under the optimal contract, such as overinvestment at time-varying degrees. We build a dynamic agency model that involves two stages. In the first—search—stage, a principal (e.g., headquarters or representative investor; referred to as "she") contracts with an agent (e.g., SPAC sponsor, VC general partner/associate, hedge fund manager, CEO, CTO, COO, or divisional manager; referred to as "he") to find investment opportunities ("projects" or "targets") that arrive randomly over time. The search requires spending of resources by the principal, but the arrival and the quality of the target are privately observed by the agent. The principal decides whether to seize the investment opportunity (and begin production) or to terminate the search process. In the second—production—stage, the agent generates an output that depends on the project quality. The principal's objective is to optimally design the contract that maximizes her total return net of the agent's compensation and provides incentives for the agent to report his private information truthfully.

Absent information frictions, the principal's first-best strategy is to invest in the first target that clears a constant and sufficiently high bar of quality (threshold) and never abandon the search. The information frictions in the second-best scenario, however, creates an adverse selection problem that distorts the principal's investment and production strategies. In the production stage, the optimal contract takes the form of an output-sharing rule, under which the agent keeps a fraction of the output as his compensation. To provide incentives for truthful reporting of the target quality, the agent's compensation must exceed his cost of production effort in the form of information rent. Under the optimal contract the equilibrium output increases with target quality, implying more information rents for better targets.

In the search stage, the optimal contract takes the form of a budget for search resources combined with—as in the first-best—a threshold for investment. The principal transfers resources to the agent continuously and subtracts them from the budget at a specified rate ("charge"). Intuitively, this structure of the optimal contract arises because only the agent observes the arrival and quality of the target. Since his information rent increases with target quality, the agent is always tempted to conceal the arrived target and wait for a better one in the future. To provide the agent with the incentives to make not only accurate but also timely report, the optimal contract grants the agent a finite budget with a progressively declining balance and a threshold for investment. If the agent announces a target with quality that clears the investment threshold, production takes place, and the agent is paid according to the terms of the contract. Otherwise, the contract is terminated without payment to the agent when the budget is exhausted. Thus, waiting becomes costly to the agent as it increases the likelihood that no target that clears the investment threshold arrives before his contract is terminated without pay.

Unlike the first-best scenario and the predictions of static models (e.g., Bernardo, Cai, and Luo, 2001; Baldenius, 2003; Heinle, Ross, and Saouma, 2014; Bastian-Johnson, Pfeiffer, and Schneider, 2017), the investment threshold, the internal charge for search resources, and the balance of the agent's budget are all endogenous and time-varying as a result of the search history. Crucially, while the *threat* of termination is a necessary incentive device due to the agent's private information, it is also costly to the principal whose payoff comes solely from the output of production. Therefore, conditional on maintaining proper incentives, the investment threshold under the optimal contract reflects a balance between the benefit of continuing the search (the prospect of high-quality targets arriving in the future) and the cost of doing so (the increasing likelihood of termination when the budget is exhausted). Early on in the search, the investment threshold is high because the likelihood of contract termination is not too great of a concern. As time elapses and in the absence of a suitable target, the balance of the resource budget declines while the likelihood of termination increases. Thus, the investment hurdle decreases over time to expedite the transition to production. Eventually, when the balance of the resource budget becomes too low and termination becomes imminent, the investment threshold reaches its minimal level such that all targets regardless of their quality trigger the investment. Because such search cost is always present and progressively higher, the optimal investment threshold induced by the agency frictions is always below the first-best level and progressively decreases, resulting in *overinvestment at time-varying degrees.*

We also extend the model to incorporate two additional features that arise naturally in practice. First, we allow the agent to divert the search resources to generate private benefits, in which case no target would arrive. This is equivalent to allowing the agent to shirk during search, generating a moral hazard problem that interacts with the adverse selection problem in the baseline model. To provide the incentives for search effort, the agent's continuation utility must decline at a minimal rate commensurate to his private benefit from shirking. Consequently, this limits the maximal level of quality the principal optimally chooses to invest in, exacerbating overinvestment. Our second extension introduces dynamics in the production stage. We assume the target quality determines only the initial productivity, which afterwards evolves over time subject to time-varying shocks. The evolution path of the productivity is privately observed by the agent and the principal dynamically adjusts her production policies based on the agent's reported productivity. We find that this extension retains the main theoretical implications of the baseline model while delivering several additional insights.

Our model generates empirically testable predictions regarding special purpose acquisition companies (SPAC), mergers and acquisitions (M&A), venture capitalists (VC), hedge fund activism (HFA), internal innovation, and project selection. There is substantial empirical evidence for overinvestment in these markets, which is often interpreted as an indication of the agent's empire-building preference and a lack of internal discipline.[1] This paper offers an alternative explanation based on optimal contracting under agency frictions and develops testable hypotheses that link overinvestment to observable firm and industry characteristics. We expect that the degree of overinvestment is positively correlated with the number of firms and frequency of activities in these markets, the incentive power of executive contracts, as well as the geographical proximity, executive connections, stock liquidity, analyst coverage, and institutional holdings of the targets.

2 Related Literature

Our study is related to the literature studying investment decisions in principal-agent frameworks (e.g., Bernardo, Cai, and Luo, 2001; Baldenius, 2003; Baldenius, Dutta, and Reichelstein, 2007; Dutta and Fan, 2009; Heinle, Ross, and Saouma, 2014; Bastian-Johnson,

[1]We review the empirical evidence on overinvestment and discuss its interpretations in Section 5.

Pfeiffer, and Schneider, 2013, 2017). The agent has private information that is related to the quality or return on investment, and the optimal contract provides incentives for the truthful reporting of such information. In the absence of other frictions, this typically results in capital rationing, or underinvestment compared to the first-best. Several studies (e.g., Antle and Fellingham, 1990; Fellingham and Young, 1990; Arya, Fellingham, and Young, 1994; Baiman, Heinle, and Saouma, 2013) consider settings with multiple stages where the investment decisions of later stages depend on the agent's report in the previous rounds. However, the number of rounds is pre-determined, and the set of available investment opportunities is given in each round. Our work contributes to this literature by incorporating dynamic search and random arrival of investment targets of unknown quality, thus introducing several novel features compared to the existing studies on this topic. Most importantly, the length of the search period is stochastic and endogenous, determined by the search history. The principal must provide incentives for timely and truthful reporting of the arriving targets and their quality. The resulting optimal contract features an investment hurdle below the first-best that progressively decreases as the search continues. That is, our model predicts overinvestment at time-varying degrees.

Several recent studies incorporate dynamics in agency-based investment models with different features and implications. For example, Malenko (2019) explores the optimal design of a dynamic capital allocation process in which an agent (division manager) privately observes the arrival and quality of investment projects. Different from us, each project is a take-it-or-leave-it opportunity with an instant return if the headquarters undertake it, and there is an auditing technology (for the most part) that perfectly reveals the quality of the project at a cost. The agency friction arises from the agent's empire-building preference: he is inclined to exaggerate the quality of the project in order to induce a larger investment. As a result, his continuation utility drifts *upward* in the absence of any reported project and jumps *downward* when an investment project is taken without auditing to cancel out his private benefit from the investment. All projects, regardless of their quality, receive investment, and there is no termination. However, the principal controls the scale of each investment, with larger projects subject to more severe agency frictions in mind. Consequently, the optimal contract always induces underinvestment (in terms of scale) relative to the first-best. In contrast, in our model the agent's continuation utility drifts *downward* in the absence of a suitable target and jumps *upward* when investment is made, and the search must be terminated within finite time to provide the agent with the incentives for timely reporting. Additionally, only one target can be taken, and the principal controls the threshold of quality for such investment. Consequently, the optimal contract always induces overinvestment (in terms of the scope of quality) and yields different empirical implications.

Varas (2018) studies a dynamic model of short-termism in which the agent's private information is binary: the agent can either spend effort and time to discover a value-enhancing good project, or pass off a value-destroying bad project which is always available. The principal uses a contract with decreasing compensation as the incentive for effort. However, to prevent short-termism, the contract holds the compensation stationary at some point and switches the incentive to random termination. In the equilibrium, the bad project is never invested, i.e., short-termism never occurs. Our model differs in that there is a continuum of private information: the arrival and quality of the investment target is observable only to the agent. Thus, the optimal contract provides time-varying incentives for the truthful reporting of his private information at different points in time. As a result, the principal endogenously becomes more "short-termist" and progressively permits investment in lower-quality targets over time. Also, while the contractual relationship in Varas (2018) continues after the investment project is chosen, the agent takes no action in that period and cannot influence the project output. In contrast, the agent in our model is also in charge of production from the chosen project and can manipulate the project output to extract more rent.

More broadly speaking, our model is also related to the literature of dynamic contracting models with Poisson jumps (e.g., Biais, Mariotti, Rochet, and Villeneuve, 2010; Hoffmann and Pfeil, 2010; Piskorski and Tchistyi, 2010; DeMarzo, Fishman, He, and Wang, 2012). The most closely related studies include Green and Taylor (2016), Curello and Sinander (2021), Madsen (2022), and Mayer (2022), in which the agent can, through effort, observe a private signal that is valuable to the principal. The optimal contract provides incentives for the agent to exert the effort to uncover the signal and to report it as soon as it arrives. While we consider an extension of our baseline model that includes a similar moral hazard problem on search effort in Section 6.1, our paper differs from these studies in two dimensions. First, in addition to the arrival time, the agent also privately observes the quality of the private information and must be incentivized to truthfully convey both the arrival and the quality to the principal. Second, the contracting relationship does not end with the disclosure of the private information. The value of the information manifests in a production process, during which the agent can continue to utilize his private information to extract rents.

3 Economic Setting and Benchmark

Below we describe the model ingredients (in Section 3.1), discuss some of our assumptions (in Section 3.2) and present the first-best benchmark (in Section 3.3).

3.1 Model Description

We consider a principal (e.g., firm headquarters or a representative owner/investor; henceforth referred to as "she") contracting with an agent (e.g., SPAC sponsor, VC general partner/associate, hedge fund manager, CEO, CTO, COO, or divisional manager; henceforth referred to as "he"). The principal has deep pockets while the agent is protected by limited liability. Both parties are risk-neutral with no discounting, and their outside options are normalized to 0.

There are two stages—search and production—and time is continuous. During the search stage, the principal pays a flow cost δ (e.g., working capital), and targets arrive, one at a time, via a Poisson jump process N_t with intensity λ. This could represent a CTO searching to adopt a new technology, an activist hedge fund manager seeking an undervalued firm for intervention, a VC general partner looking for a promising startup, or a SPAC sponsor hunting for a valuable private firm to merge with. Each target is characterized by its quality θ, which follows a Pareto distribution with cumulative distribution function $F(\theta) = 1 - \left(\frac{\theta_{\min}}{\theta}\right)^{\kappa}$ and probability density function $f(\theta) = \frac{\kappa\theta_{\min}^{\kappa}}{\theta^{\kappa+1}}$. The distribution scale parameter is positive, $\theta_{\min} > 0$, and the shape parameter is sufficiently large, $\kappa > 2$.[2] Importantly, only the agent observes the arrival of targets and their quality θ. This can be interpreted as only the agent possesses the skills and expertise to recognize an investment opportunity at the

[2] These are technical assumptions that we discuss in Section 3.2.

moment it is available and assess its value. The principal decides whether to invest in the target based on information reported by the agent. The length of this stage is endogenous: the search ends either when the principal invests and moves on to the next (production) stage, or when the contract is terminated.

During the production stage, the agent generates output from the target chosen by the principal during the previous stage. We use the term "production" loosely to refer to any utilization of the discovered target (e.g., implementation of the innovative technology, operation of the merged firms, revamping of the target company, etc.). The production technology is $y(\theta, e) = \theta e$, where e is the agent's unobservable production effort exerted at a quadratic personal cost $h(e) = e^2/2$. Similar to the first stage, the target quality is privately observed only by the agent. The output y, however, is observable by the principal and can therefore be contracted on.

The main agency friction of the model is adverse selection stemming from the agent's private information about the arrival and the quality of each target. In particular, the unobservable effort in the production stage does not introduce an independent moral hazard problem. It only provides cover for the agent so that the true quality of the target cannot be inferred with certainty based on the observable output.[3]

In line with prior literature, a contract $\mathcal{C}$ between the principal and the agent consists of the principal's investment and production policies, and the associated compensation to the agent. Even though the contract is signed at the onset of the game it can be separately defined in terms of incentives related to the search and incentives related to the production. The part related to the search stage specifies the set of targets that will be invested in, the reward for the agent for announcing the arrival of the target, and the condition under which the contract is terminated. The part related to the production stage specifies how the agent will be compensated based on the observed production. A contract is incentive-compatible if the agent finds it optimal to always announce his private information truthfully.

[3]In other words, although the production stage features both hidden information and hidden action, the agent only has one degree of freedom in his choices, a point that will become clearer as we present the solution of the model in Section 4.1.

3.2 Discussion of Assumptions

We now discuss in detail several simplifying assumptions of the model and explain why they are *not crucial* for our analysis.

1. *Discounting.* Our model assumes that there is no discounting for the principal and the agent. This assumption is common among models in which the arrival of information follows a Poisson process (e.g., Green and Taylor, 2016; Mayer, 2022). Discounting implies that the principal prefers earlier resolutions, which further distorts her investment threshold downward (beyond the one predicted by our model). Except for this result, discounting usually adds little economic intuition in these types of models but a substantial degree of algebraic complexity.

2. *Continuous time.* The assumption that time is continuous allows for a more elegant analysis. A setting with discrete time leads to several analytical complications (such as the need for randomized termination) but does not qualitatively change our results.

3. *Pareto distribution.* We assume that the target quality θ follows a Pareto distribution because of its broad applications in economics and its analytical convenience. In particular, this distribution belongs to the power-law family and is descriptive of many economic variables and activities in practice (e.g, Gabaix and Landier, 2008). Furthermore, the distribution has two analytical advantages. First, the inverse hazard rate $[1 - F(\theta)]/f(\theta) = \theta/\kappa$ is a linear function of θ. This significantly simplifies the proof of Proposition 1 and all subsequent analyses. Second, a Pareto distribution truncated from below at an arbitrary point $x > \theta_{\min}$ is also a Pareto distribution with the same shape parameter and the new scale parameter x. This is vital for tractability. The requirement $\kappa > 2$ is merely technical and implies a sufficiently thin right tail of the distribution to ensure a finite variance of θ and a finite solution to the first-best (see footnote 5 for more details).

4. *Production technology.* Our model assumes that the production technology is $y(e, \theta) = e\theta$. This implies that effort and project quality are perfect complements and achieves two useful simplifications. First, in equilibrium, production effort is never shut down

regardless of target quality.[4] Second, while the agent can misreport the investment quality, he cannot generate output without a project (e.g., from a phony acquisition target or without adopting a new technology).

In summary, the assumptions discussed above facilitate tractability and are not crucial for our results. The predictions of the model remain qualitatively unchanged if we relax any of these assumptions (e.g., switch to discrete time, allow discounting, consider a dynamic production process, use an alternative distribution for θ and/or production function for y).

3.3 First-Best Benchmark

When all information is public, the first-best effort and output in the production stage maximize the social surplus from production. i.e.,

$$\max_e \; y - h(e) = \theta e - h(e). \tag{1}$$

The first-best effort and output,

$$e^{FB} = \theta, \tag{2}$$

$$y^{FB} = \theta^2, \tag{3}$$

are both increasing functions in θ. The agent is only compensated for his cost of effort, and the payoff to the principal is $V_2^{FB}(\theta) = \theta^2/2$.

The search stage under the first-best scenario represents a standard bandit problem. Let Θ^{FB} denote the set of targets that will be invested in and V_1^{FB} denote the principal's expected value at the outset of the search stage. The following lemma summarizes the first-best investment policy:

Lemma 1 *Under the first-best scenario,* $\Theta^{FB} = \{\theta : \theta \geq x^{FB}\}$*, where*

$$x^{FB} = \left[\frac{\lambda \theta_{\min}^{\kappa}}{\delta(\kappa - 2)}\right]^{\frac{1}{\kappa-2}}. \tag{4}$$

[4] In contrast, if the production technology is linear, e.g., $y = e + \theta$, then under the optimal contract, effort may be shut down if target quality θ is sufficiently high.

The principal's first best expected payoff at the beginning of the search stage is

$$V_1^{FB}(x^{FB}) = \frac{\kappa\left(x^{FB}\right)^2}{2(\kappa-2)} - \frac{\delta\left(x^{FB}\right)^{\kappa}}{\lambda\theta_{\min}^{\kappa}}. \tag{5}$$

The optimal strategy of the principal under the first-best scenario is to finance the search with a constant minimal (cutoff) quality x^{FB}.[5] Because x^{FB} also determines the expected duration of the search, Lemma 1 suggests that, on average, the principal expects to wait longer in the first-best scenario if there are many opportunities on the market (high λ) or if searching is cheap (search flow cost δ is low).

4 Optimal Contract Under Asymmetric Information

We now analyze the optimal contract when the arrival and quality of the targets are the agent's private information. To do so, we consider separately (in backward order) the terms that provide incentives to the agent in each of the stages.

4.1 Production Stage

As a starting point, we consider the following reduced problem without the search stage: the agent is endowed with a target of quality θ that is unobservable by the principal. The agent has reservation utility W_{τ^-}, which in this reduced problem is given (but in the full-fledged problem represents the utility carried over from the search stage). The principal must design a screening contract that solicits truthful reporting of θ during the production stage while maximizing her payoff, which is the output net of the agent's compensation.

Based on the Revelation Principle, we can—without loss of generality—consider the screening contract as a direct mechanism: the agent reports his type $\widehat{\theta}$ and receives an output target $y(\widehat{\theta})$ and associated compensation $w(\widehat{\theta})$ if and only if the output target is produced.[6] Given the contract, the agent's objective is to maximize his compensation net of

[5]Equations (4) and (5) illustrate the need to assume $\kappa > 2$. Otherwise, x^{FB} and V_1^{FB} are not well-defined.

[6]Such contract is feasible because, given θ, there is no uncertainty or noise in the production technology. Note that this contract can be alternatively written in a standard "pay-for-performance" form: a function $w(y)$, under which the agent is free to produce any level of output y and receive the corresponding wage $w(y)$, and no reporting is necessary. These two formulations are equivalent because in the equilibrium, both

his (production) effort cost:

$$R(\theta) = \max_{\widehat{\theta}} \quad w(\widehat{\theta}) - h(e) \tag{6}$$

subject to the constraint

$$e = y(\widehat{\theta})/\theta, \tag{7}$$

because he needs to exert the necessary effort to produce the required level of output $y(\widehat{\theta})$ in order to receive compensation. This constraint illustrates that although effort is unobservable in the production stage, the main underlying agency friction is adverse selection: effort merely provides a cover for the agent's report $\widehat{\theta}$ so that his true type θ remains hidden.

The contract is incentive compatible if and only if

$$\theta = \arg\max_{\widehat{\theta}} \quad w(\widehat{\theta}) - h\left(\frac{y(\widehat{\theta})}{\theta}\right). \tag{8}$$

When (8) is satisfied, $R(\theta)$ is known as the agent's *information rent*: the amount of utility (in excess of his reservation utility) that he must receive in order to truthfully reveal his private information.

Because the principal does not observe θ, her objective is to maximize the expected output net of the agent's compensation. The expectation is taken over the distribution of θ, which in equilibrium is the result of the investment policy to which the principal had committed before the search and production stages. Section 3.3 showed that in the first-best scenario this investment policy is a threshold one. Our next result establishes that the investment policy under asymmetric information is *also a threshold*; we denote it by x.[7]

Lemma 2 *Under asymmetric information, it is not optimal for the principal to adopt an investment policy that is not a threshold x defined as $\Theta = \{\theta : \theta \geq x\}$.*

the output target $y(\widehat{\theta})$ and the wage $w(\widehat{\theta})$ are strictly increasing in the agent's reported type $\widehat{\theta}$, thus creating a one-to-one mapping between output and wage. Remark 1 below demonstrates such equivalence in detail. Following the accepted standard in the literature, we consider the direct mechanism that involves reporting because of its transparency in demonstrating the incentive power of the contract.

[7]This result is more general and does not depend on the presence of a search stage prior to production.

Intuitively, if a target of quality $\bar{\theta}$ will trigger investment under some incentive-compatible contract, the principal can always induce truthful reporting and thus invest in all targets with better quality (i.e., $\theta > \bar{\theta}$) by setting the output target to be $y(\bar{\theta})$ and the wage to be $w(\bar{\theta})$ for those targets.[8] Therefore, the principal's expected net return must be at least weakly increasing in target quality, and excluding high-quality targets is sub-optimal. As a result, the principal's maximal payoff from the screening contract can be written as $V_2(x) - W_{\tau^-}$, where $V_2(x)$ solves

$$V_2(x) = \max_{y(\widehat{\theta}),w(\widehat{\theta})} \int_x^{+\infty} [y(\theta) - w(\theta)] \left(\frac{\kappa x^\kappa}{\theta^{\kappa+1}}\right) d\theta \tag{9}$$

$$= \max_{y(\widehat{\theta}),w(\widehat{\theta})} \int_x^{+\infty} [y(\theta) - h(e(\theta)) - R(\theta)] \left(\frac{\kappa x^\kappa}{\theta^{\kappa+1}}\right) d\theta \tag{10}$$

subject to the IC condition (8).[9] In other words, given the investment policy x, $V_2(x)$ captures the principal's expected payoff from production under the incentive-compatible screening contract with optimally designed output-compensation combinations.

Deriving $R(\theta)$ and $V_2(x)$ represents a static mechanism design problem to which the solution is summarized as follows:

Proposition 1 *Let $\gamma \equiv \frac{\kappa}{\kappa+2} < 1$. Given any investment threshold $x \geq \theta_{\min}$, the optimal screening contract in the production stage has the following properties:*

- *The agent's information rent from a target of quality $\theta \geq x$ is given by*

$$R(\theta) = \frac{\gamma^2}{2}(\theta^2 - x^2); \tag{11}$$

- *The principal's expected payoff from production is given by*

$$V_2(x) = \frac{\gamma\kappa}{2(\kappa-2)} x^2; \tag{12}$$

- *The optimal output is $y^* = \gamma\theta^2$ and the production effort is $e^* = \gamma\theta$.*

[8] This is incentive compatible because it requires less effort from the agent to produce $y(\bar{\theta})$ when $\theta > \bar{\theta}$.

[9] The term $\frac{\kappa x^\kappa}{\theta^{\kappa+1}}$ represents the distribution of θ given the investment threshold x.

The agent's information rent is a quadratic function of target quality θ and the principal's expected payoff is a quadratic function of the investment threshold x. In particular, (11) implies that the *marginal* rent $R'(\theta) = \gamma^2\theta$ is a linearly increasing function of target quality θ. This is known as the *slope condition* which is essential in the presence of adverse selection for ensuring incentive compatibility, that is, the truthful reporting of θ. Meanwhile, compared with the first-best level of effort e^{FB} in (2) and the first-best level of output y^{FB} in (3), adverse selection distorts both the optimal effort e^* and output y^* downward by a constant fraction: $1 - \gamma = 2/(\kappa + 2)$. Intuitively, a higher κ corresponds to a lower variance in θ and therefore, less information asymmetry.

Given x, equation (11) implies that the conditional expectation of the information rent the agent can receive is

$$U(x) \equiv \mathrm{E}\left[R(\theta)|\theta \geq x\right] = \int_x^{+\infty} \frac{\gamma^2}{2}\left(\theta^2 - x^2\right)\left(\frac{\kappa x^\kappa}{\theta^{\kappa+1}}\right) d\theta = \frac{\gamma^2 x^2}{\kappa - 2}. \tag{13}$$

The closed-form expressions for $U(x)$ and $V_2(x)$ greatly simplify the design of the optimal contract in the search stage in the next section.

Remark 1 *The screening contract in Proposition 1 can be implemented via a simple output sharing rule*

$$w(y) = \gamma(y - \gamma x^2) + \frac{\gamma^2 x^2}{2}. \tag{14}$$

Under this rule, if the agent produces a minimal amount of output γx^2, he receives a basic wage $\frac{\gamma^2 x^2}{2}$ which exactly offsets his effort cost. Then, for every additional unit of output the agent produces, he receives γ fraction of that as his compensation. This simple output sharing rule represents an indirect mechanism, under which the agent does not need to report his type and can freely produce any level of output he desires. In comparison, Proposition 1 is derived based on a direct mechanism, under which the agent reports $\widehat{\theta}$ and receives an output target $y(\widehat{\theta})$ and the corresponding wage $w(\widehat{\theta})$. However, these two mechanisms are equivalent (as a result of the revelation principle), because they are both incentive compatible and deliver the exact same rent $R(\theta)$ and the same expected payoff to the principal. Thus, our choice of *formulating the solution to the adverse selection problem in the production stage as a direct mechanism is without loss of generality.*

4.2 Search Stage

The analysis in the previous section solves the optimal screening contract that induces truthful reporting of target quality from the agent, if the existence of the target is known. During the search stage, however, the arrival of a target is random and only observable to the agent. Moreover, Proposition 1 shows that $R(\theta)$, the information rent for truthful reporting, is an increasing function, meaning that the agent receives higher rent from a target of higher quality. This gives the agent the incentive to conceal the arrival of a target from the principal and wait for a better one in the future. Since the support of θ is unbounded, the agent always prefers to wait. To deter this, the agent must be given the proper incentives to report the arriving targets both truthfully and *in time*. This is the main objective behind the optimal contract in the search stage, which is derived below.

Similar to standard dynamic agency models (especially those set in continuous time such as DeMarzo and Sannikov, 2006; Biais, Mariotti, Plantin, and Rochet, 2007; Sannikov, 2008), we find that incentives for timely reporting in our model can be provided in the form of promised future compensation to the agent. Specifically, let $\{C_t\}_{t\in[0,\tau]}$ denote the compensation to the agent during the search stage, τ denote the stopping time either because of transition to production stage or contract termination, and W_τ denote the terminal payment. The contract can be characterized using the agent's continuation utility W_t, defined as

$$W_t = \mathrm{E}\left[\int_t^\tau dC_s + W_\tau\right]. \tag{15}$$

We can solve the optimal contract that provides incentives to the agent during the search stage as follows. First, because all players are equally patient, any intermediate compensation can always be delayed at no cost. Given that the production stage is static without noise or risk, it is without loss of generality to accrue all payments until the output is produced. That is, under the optimal contract, the agent is paid if and only if the contract moves to the production stage and the agent produces the required output (i.e., $dC_t = 0$ for all $t < \tau$

and $W_\tau = 0$ if the contract is terminated without production). The following proposition thus characterizes the dynamics of the agent's continuation utility in the search stage:

Proposition 2 *Let x_t denote the investment threshold set by the optimal contract at time t, the agent's continuation utility W_t evolves according to $dW_t = U(x_t)dN_t - J(x_t)dt$, where*

$$J(x_t) = \left[\lambda \left(\frac{\theta_{\min}}{x_t}\right)^\kappa\right] U(x_t) = \left(\frac{\lambda\gamma^2\theta_{\min}^\kappa}{\kappa - 2}\right) x_t^{2-\kappa}. \tag{16}$$

The contract is terminated if $W_t = 0$.

Similar to standard models with Poisson search, W_t jumps upward and the firm moves on to the production stage if a suitable target arrives. Here, a target is suitable if its quality clears the threshold of investment, i.e., $\theta_t > x_t$. Otherwise, W_t drifts down at rate $J(x_t)$ if no target arrives or is reported. If a sufficiently long time has passed without reporting, the search terminates at $W_t = 0$, and the agent receives no payment. Put differently, to dissuade the agent from waiting for a better target, the agent must be worse off waiting. The contract achieves this using the *threat* of termination without pay. Such threat becomes more imminent, that is, the likelihood of termination increases, as time passes by without a report of a suitable target. Consequently, W_t drifts down over time and jumps up only if a target of sufficient quality is reported.[10] However, different from the standard models, the adverse selection in the production stage in our model implies that W_t jumps up by a specific value dictated by the solution to the optimal screening contract (i.e., Proposition 1). Given the investment policy x_t, $U(x_t)$ is the agent's expected utility reward (i.e., the expected size of the upward jump in W_t) if a suitable investment target arrives. Multiplying that by the probability that a suitable target arrives, $\lambda \left(\frac{\theta_{\min}}{x_t}\right)^\kappa$, yields $J(x_t)$ in (16).

The dynamics of the contract in Proposition 2 can be implemented by assigning the agent a "budget" for the spending of search resources. The budget has a balance W_t, and decreases at rate $J(x_t)dt$ until either the agent reports the arrival of a suitable target or the budget is

[10] W_t does not jump if the agent reports a target that falls below the quality threshold x_t. Such target does not trigger investment, and thus the agent's report is not verifiable. Moreover, an upward jump in W_t implies a faster decline of W_t, which increases the likelihood of contract termination and is costly to principal. Therefore, it is optimal for the principal to make payments only when production occurs.

depleted. The former triggers production and compensation as described in Proposition 1. The latter triggers termination of the search process.

The principal has two controls when designing the optimal contract: the investment threshold x_t, and the terminal compensation W_τ. Her expected payoff at any time $t \in [0, \tau]$ under the optimal contract, denoted as $V_{1,t}$, solves

$$V_{1,t} = \mathrm{E}\left[\int_t^\tau -\delta ds + y_\tau - W_\tau\right], \tag{17}$$

subject to the IC constraints (8). It holds that $y_\tau = y$ if production takes place, and $y_\tau = 0$ if the contract is terminated without production. The principal pays for the search cost δ. She retains the production output y if investment is made, but has to pay terminal compensation W_τ to the agent. The analysis so far has pinned down the optimal, incentive compatible terminal compensation: if a suitable target can be found, the agent receives an extra reward, which is $U(x)$ in expectation. Otherwise, if the contract terminates without production, he receives no payment. Thus, the ensuing analysis focuses on characterizing the optimal investment policy x_t.

Proposition 2 implies that the principal's payoff under the optimal contract can be summarized as a function of the agent's continuation utility, or $V_1(W)$, which solves the following Hamilton-Jacobi-Bellman (HJB) equation with x being the only control variable:

$$0 = \max_x \quad -\delta - \lambda\left(\frac{\theta_{\min}}{x}\right)^\kappa U(x)V_1'(W) + \lambda\left(\frac{\theta_{\min}}{x}\right)^\kappa [V_2(x) - W - V_1(W)]. \tag{18}$$

The first term represents the search cost. The second term stems from the drift of dW_t, and the third term represents the change in the principal's payoff if a suitable target is found and the contract moves into production.[11]

[11] Conditional on moving into production, the principal's final payoff is the output y net of the wage $w(\theta)$ paid for production and the agent's residual utility W carried over from the search stage. The former can be further divided into the compensation for the agent's production effort $h(e)$ and his information rent $R(\theta)$, all embedded in the definition of $V_2(x)$ in equation (10). Put differently, the principal's final payoff can be written as $y - R - h(e) - W$, where the first three terms are captured (in expectation) by V_2.

Rearranging terms, the HJB equation can be conveniently written as

$$V_1(W) = \max_x \quad V_2(x) - W - U(x)V_1'(W) - \frac{\delta}{\lambda}\left(\frac{x}{\theta_{\min}}\right)^{\kappa}. \tag{19}$$

The tradeoff faced by the principal when setting the optimal investment threshold is as follows: A higher x yields a higher expected payoff once a suitable target arrives: $V_2'(x) > 0$, as seen in (12). The cost, however, is two-fold. First, targets with high quality arrive at a lower rate which leads to higher search cost in expectation, the last term in (19). Second, once a target arrives, the agent is given a higher reward to truthfully reveal the target quality: $U'(x) > 0$, as seen in (13). This higher reward must be accompanied by a faster decline of W to maintain W as a martingale, which increases the likelihood of contract termination.

The optimal choice of threshold x can be obtained by the first-order condition:

$$V_2'(x) - U'(x)V_1'(W) - \left(\frac{d}{dx}\right)\left[\frac{\delta}{\lambda}\left(\frac{x}{\theta_{\min}}\right)^{\kappa}\right] = 0. \tag{20}$$

The first term is positive, because x increases the expected payoff for the principal in the production stage. The third term is negative, because x increases the expected time to wait and thus the expected search cost. The middle term captures the marginal continuation utility and its sign depends on W. When W is large, $V_1'(W) < 0$, because the promised compensation to the agent lowers the principal's payoff if a suitable target arrives and the search ends. When W is small, $V_1'(W) > 0$, because the primary concern for the principal is the likelihood of contract termination, which increases as W declines. Substituting $V_2(x)$ from (12) and $U(x)$ from (13) into (20) yields the optimal investment policy:

$$x(W) = \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right)V_1'(W)\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}. \tag{21}$$

The optimal investment threshold is increasing in the agent's continuation utility, $x'(W) > 0$.[12] Because W_t drifts down over time, the optimal investment strategy is to adopt a progressively lower threshold for the quality of the arriving targets worth investing in.

[12] Technically speaking, this is because $V_1(W)$ is a concave function, i.e., $V_1''(W) < 0$, which is a standard feature of dynamic contracting models and is visualized in Figure 1.

Meanwhile, $\theta_{\min}$ represents the lowest investment threshold that the principal can set. Substituting $x = \theta_{\min}$ into (21) implies that $x(W) = \theta_{\min}$ for all $W \leq \underline{W}$, where $\underline{W}$ solves

$$V_1'(\underline{W}) = \left(1 - \frac{\delta(\kappa-2)}{\gamma\lambda\theta_{\min}^2}\right)\frac{\kappa}{2\gamma}. \tag{22}$$

That is, when W is sufficiently low, the optimal policy is to invest in the next target that arrives, regardless of its quality. This maximizes the probability that the contract moves to the production stage before it is terminated.[13]

Altogether, the optimal contract during the search stage can be summarized in the following proposition:

Proposition 3 *Under the optimal contract, the principal's value function $V_1(W)$ solves the HJB equation (18) subject to the boundary condition $V_1(0) = 0$. The optimal investment policy $x(W)$ is given by*

$$x(W) = \begin{cases} \theta_{\min}, & \text{if } W \leq \underline{W} \\ \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right)V_1'(W)\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}, & \text{if } W > \underline{W} \end{cases} \tag{23}$$

where $\underline{W}$ is given by (22). $x'(W) > 0$ for all $W > \underline{W}$.

Figure 1 illustrates the value function $V_1(W)$ and the optimal investment policy $x(W)$. Without any further constraint on the agent's initial outside options, W_t optimally starts at $W_0 = W^*$ where $V_1(W)$ is maximized, and drifts to the left until either production or termination occurs. Figure 1 also plots the first-best investment policy x^{FB} and demonstrates the following result:

[13] The result that all targets trigger production when W is sufficiently low relies partially on the assumptions that all targets, regardless of their quality, generate positive returns to the principal and require time to be discovered. The former can be justified if the principal has a common sense of the basic properties of investment opportunities worth taking (e.g., firms with strong growth history and healthy balance sheet), and the latter can be interpreted as "no free lunch" in the financial market. If, instead, the $\theta_{\min}$ target represents a "default" option that is always immediately available, then when W is sufficiently low, the principal will intuitively abandon the search by resorting to the default option in lieu of contract termination. If $\theta_{\min} < 0$, or if there is a substantial fixed cost for production, then the optimal investment policy may exclude some low-quality targets even when W is low and termination is imminent. If $\theta_{\min}$ is always available and low-value, the optimal contract may involve random termination in order to peg W at a sufficiently high level to prevent the agent from exploiting this low-value default option, such as the case studied in Varas (2018).

Corollary 1 *It holds that $x(W) < x^{FB}$ for all W.*

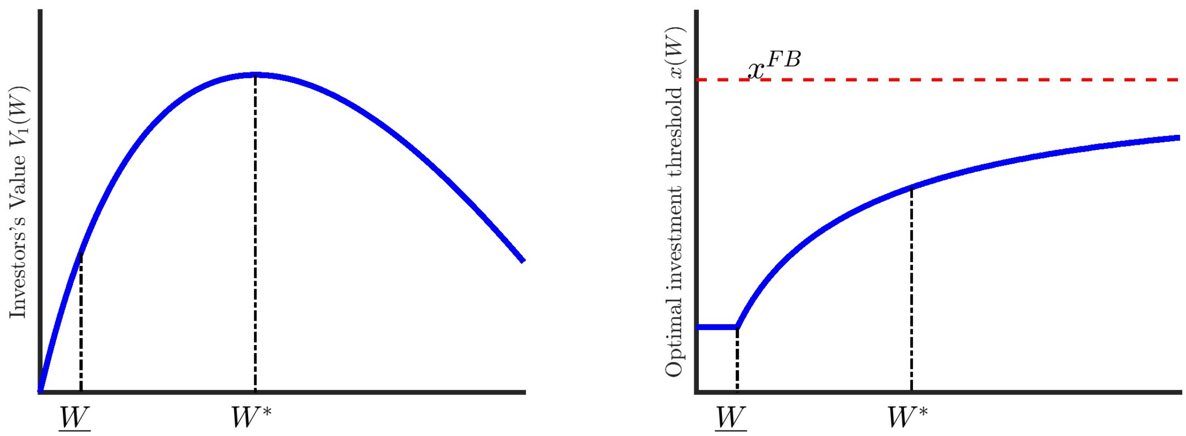

Figure 1: **The Principal's Value Function and The Optimal Investment Policy**
The left panel of this figure plots the principal's value function $V_1(W)$ under the optimal contract. The right panel plots the optimal investment threshold $x(W)$ according to Proposition 3 and the first-best investment threshold x^{FB}. $\underline{W}$ is defined according to (22), and $W^* \equiv \arg\max_W V_1(W)$ represents the point at which $V_1(W)$ is maximized. Parameter values are $\lambda = 2.5$, $\kappa = 4.25$, $\delta = 1.1$.

Combining the observations that $x(W)$ is increasing in W and the latter drifts over time, we can immediately show:

Corollary 2 *The optimal investment policy $x(W)$ is progressively lower over time.*

To reduce the likelihood of contract termination, the principal reduces the investment threshold. When W is sufficiently low, all targets trigger production regardless of their quality. When W is larger, the concern for termination is somewhat eased but is never eliminated because there is a strictly positive probability of contract termination regardless of how large W is. Hence, the optimal investment threshold in the presence of agency frictions is always below the first best, i.e., firms overinvest. The threshold increases in W, which in turn decreases in time. Thus, overinvestment becomes progressively more severe.

Notably, a progressively decreasing investment threshold x also implies that the speed at which the agent's budget balance decreases, $J(x)$, is accelerating over time. One way the principal can control the agent's budget balance is to internally charge him $p = J(x)/\delta$ for the resources spent in the search process. A higher $J(x)$ implies a higher p. Since $J(x)$, given in (16), is a decreasing function of x, Corollary 2 implies the following result:

Corollary 3 *The internal charge for search resources, $p(W)$, is progressively higher over time.*

The charge in Corollary 3 can exceed one for a large set of parameter values. In this case, we predict that firms may impose a charge that is higher than their marginal cost, a result that is consistent with Baldenius, Dutta, and Reichelstein (2007).

In summary, the optimal contract providing incentives for the search stage can be implemented via a resource budget combined with a time-varying investment threshold/hurdle. Resources spent for search are subtracted from that budget at a (time-varying) rate commensurate with the search history. Early on in the search, the investment hurdle is high and the charge for resources is low, implying a slow decrease of the budget balance. As time elapses and in the absence of a suitable project, the principal optimally lowers the investment hurdle to expedite the transition to production while raising the charge to maintain proper incentives for truthful and timely report of the arriving target. These dynamics are the result of balancing the benefit of search (the prospect of high-quality projects arriving in the future) with the cost of doing so (the threat of termination when the budget is exhausted).

4.3 Discussion of Investment Distortions

One of our main results is overinvestment—it differs from those in prior studies that also feature adverse selection in project choices. In this section we illustrate how such result arises using a stylized example to which our model can be related.

Consider a generic production technology characterized by the output function $y(\theta)$ and the cost function $c(\theta)$. Both are weakly increasing functions of θ, which measures the quality or type of the project. A risk-neutral principal, who bears the cost of production, contracts with a risk-neutral agent, who operates the technology and thus may have superior information about the value of θ. In such case of asymmetric information, the principal only knows that θ is drawn from a pool of projects characterized by a distribution function $F(\theta)$ with unbounded support. There are two types of project selection problems in this generic setup:

Problem I: The agent draws a single project out of the pool of projects, and the principal's decision is to invest or not. That is, the principal faces a *one-shot, take-it-or-leave-it* *decision for a given project.* Investment incurs the production cost c and yields the output y, and no-investment yields zero net payoff. This is the standard static investment model that is commonly used in prior studies with varying details (e.g., Baldenius, 2003; Baldenius, Dutta, and Reichelstein, 2007; Bastian-Johnson, Pfeiffer, and Schneider, 2013).

We begin with analyzing the first-best project choice, denoted by x^{FB}, when θ is observable. A common result is that x^{FB} is a cutoff that represents the minimal quality of project that triggers investment. That is,

$$x^{FB} = \min\{\theta : y(\theta) - c(\theta) \geq 0\}. \tag{24}$$

In other words, $y(x^{FB}) = c(x^{FB})$, and all projects with $\theta \geq x^{FB}$ are undertaken. This result holds as long as $y(\theta) - c(\theta)$ is weakly increasing in θ.

Now suppose only the agent sees the true value of θ and must be given incentives to report it truthfully to the principal. This creates a canonical adverse selection problem which, depending on the specific assumptions, yields different incentive-compatible solutions. However, a typical result of this adverse selection friction is that the equilibrium output is lower than that under the first-best. For the purpose of illustration, we assume that it reduces the principal's payoff from each project to $\gamma y(\theta)$ where $\gamma < 1$. This can be interpreted as $1-\gamma$ fraction of the output must be paid to the agent as his information rent in exchange for truthful reporting.[14] Then, the principal's optimal choice under adverse selection, denoted by x^*, solves

$$x^* = \min\{\theta : \gamma y(\theta) - c(\theta) \geq 0\}. \tag{25}$$

In other words, $\gamma y(x^*) = c(x^*)$. The question of interest is how x^* compares to x^{FB}. To answer that, define the cost-to-output ratio function $\phi(\theta) = c(\theta)/y(\theta)$. As a result, we have $\phi(x^{FB}) = 1$ and $\phi(x^*) = \gamma < 1$. The assumption that $y(\theta) - c(\theta)$ is weakly increasing in θ,

[14] In our model above, this is indeed the *optimal* screening contract (see Proposition 1 and Remark 1), with $\gamma = \kappa/(\kappa+2)$. This results from the combination of multiplicative production for y and Pareto distribution for θ. Other combinations of assumptions can result in different γ. However, the specific value of γ and whether it is a constant are not critical. The basic intuition in this section is intact as long as $\gamma(\theta) < 1$ for all θ, which is a standard result of most adverse selection models.

which is necessary for x^* and x^{FB} to represent investment cutoffs, implies that $\phi(\theta)$ must be decreasing in θ.[15] Then, we have $x^* > x^{FB}$, which implies *underinvestment*.

Problem II: Now imagine that instead of a one-shot, take-it-or-leave-it decision for a given project, the principal is *presented with the whole set of projects and selects the optimal one to investment in*. In the first-best, her optimal choice solves

$$x^{FB} = \arg\max_{\theta} \ y(\theta) - c(\theta). \tag{26}$$

When the objective function $y(\theta) - c(\theta)$ is twice differentiable and globally concave the solution is given by the first-order condition $y'(x^{FB}) = c'(x^{FB})$. Assuming once again that adverse selection reduces the output by $\gamma < 1$, the principal's optimal choice x^* now solves

$$x^* = \arg\max_{\theta} \ \gamma y(\theta) - c(\theta). \tag{27}$$

The solution satisfies $\gamma y'(x^*) = c'(x^*)$. The question of interest is still how x^* compares to x^{FB}. Analogously, we can define the *marginal* cost-to-output ratio $\psi(\theta) = c'(\theta)/y'(\theta)$ and the assumption of global concavity implies that the marginal cost must increase faster than the marginal revenue, which implies that $\psi(\theta)$ must be increasing in θ.[16] Consequently, $x^* < x^{FB}$, which implies *overinvestment*.

Comparison and discussion: In both problems, adverse selection effectively reduces the principal's payoff from a given project. In Problem I, the principal faces a single take-it-or-leave-it opportunity drawn from the pool of projects, and the optimal choice is made by finding the point where the principal's *participation constraint* ($y - c \geq 0$) binds. Consequently, we refer to this as a "participation" problem. In Problem II, the principal faces the entire set of opportunities and makes her optimal choice by *optimizing* over the benefits and the costs, which we refer to as an "optimization" problem. Using these definitions we reach the following observation:[17]

[15] A typical and widely-used example is that $y(\theta)$ is strictly increasing in θ while $c(\theta)$ is a constant (e.g., a fixed cost for investment).

[16] A typical and widely-used example is that $y(\theta)$ is a concave function and $c(\theta)$ is linear. In our model (e.g., by the HJB equation (19)), $y(\cdot) \approx (\cdot)^2$ is convex while $c(\cdot) \approx (\cdot)^\kappa$ with $\kappa > 2$, which is more convex.

[17] The analysis above is a demonstration of the basic intuition behind two types of problems commonly

Observation 1 *Adverse selection leads to underinvestment in participation problems and overinvestment in optimization problems.*

Our model yields an "optimization" problem and therefore, overinvestment. In particular, although x represents a threshold for investment, the principal's choice of x at time t is the result of an "optimization" problem indicated by the HJB equation (19). The benefit of a higher x is captured by $V_2(x)$, which is an increasing function, while the cost of a higher x is the longer waiting time.[18] The crucial mechanism at work in our model is the *dynamic* nature of the model: *by not investing in a target, the principal preserves the possibility that a better target arrives in the future.* Consequently, even though targets only arrive one at a time, the principal has the option to adjust her investment policy x through an "optimization" problem over its benefits and costs. Our model thus provides a *micro-foundation* for an "optimization"-problem-based dynamic search for investment opportunities. It is a natural characterization of the practices of internal innovation, project selection and external investment opportunities discussed in Section 1. Moreover, the dynamic nature yields a *time-varying* investment policy and therefore, time-varying degrees of overinvestment. We discuss the practical implications of these results in more details in Section 5.

5 Empirical Predictions of the Main Setting

This paper studies an investment decision in the presence of agency frictions (i.e., unobservable arrival of investment targets and their quality) combined with dynamic search for opportunities. In contrast to the underinvestment distortion found in the vast majority of previous theory research on project selection, our analysis demonstrates that the interplay of these frictions and optimal contracting results in overinvestment—a phenomenon sub-

used in models of project selection. It is not a proof for the precise conditions under which under- or overinvestment can arise. Such proof is difficult, if not impossible, unless very specific model assumptions are made, which naturally limits its generality and applicability across different models.

[18]More specifically, the cost from longer waiting can be further decomposed into the direct search cost and the indirect agency cost. In particular, a higher x requires strong incentives for the agent ($U(x)$), and the marginal cost of providing such incentives is captured by the marginal value of the agent's continuation utility ($V_1'(W)$). Because $V_1(W)$ is concave, and W decreases over time, a longer waiting time increases $V_1'(W)$ and thus the agency cost.

stantiated by empirical observations in both external and internal investment endeavors.[19] For instance, following the boom of SPACs in recent years, Dambra, Even-Tov, and George (2021), Lin, Lu, Michaely, and Qin (2021), and Gahng, Ritter, and Zhang (2022) examine the subsequent performance of SPACs and find extensive evidence of underperformance, suggesting that SPAC sponsors select low-quality private firms as targets and push those deals through. Lang, Stulz, and Walkling (1991), Harford (1999), Andrade, Mitchell, and Stafford (2001), and Masulis, Wang, and Xie (2007, 2009) document wide-spread negative stock returns for the bidders of M&A activities as a sign of those firms making value-destroying acquisition choices. In the market of hedge fund activism (HFA), Yin and Zhu (2023) find that although HFA targets in general experience positive abnormal returns around the public announcement of the hedge funds' intervention, those targets do not perform better than the other holdings of the same hedge funds and can underperform. Meanwhile, Richardson (2006) constructs an accounting-based measure and finds ample evidence for overinvestment defined as investments beyond those required to maintain assets in place and to finance expected new investments in positive NPV projects.[20] Billett, Garfinkel, and Jiang (2011) document similar findings explicitly for lumpy internal investment that result in lower operating outcomes (EBIT/assets ratio) and stock returns.

Within our model, in exchange for his private information about the quality of each arriving target, the agent receives sufficiently high information rent that increases with target quality. However, because targets arrive one at a time and the agent is the only one observing them, he is tempted to conceal the arrival and wait for a better target in the future. To deter that, the optimal contract utilizes a combination of resource budget, threshold for investment, and the threat of termination without pay. The balance of the budget decreases over time until either the agent announces a suitable target that meets the predetermined threshold (which triggers production), or the budget is exhausted (which triggers termination of the search). Waiting is now costly to the agent, as it increases the likelihood of termination

[19] Prior theoretical work often assumes that there exists only one investment opportunity—an assumption that, in the absence of additional frictions, results in underinvestment (see Section 4.3 for details). Controlling for the availability of investment opportunities (e.g., by controlling for industry), one might test the sign of investment distortions predicted in our and prior work.

[20] In Richardson (2006), investments are defined as the sum of a firm's capital expenditure (CAPEX), R&D, acquisition, and sales of PP&E.

before the next suitable target can arrive. Meanwhile, terminating the search is also costly to the principal as her payoff only comes from the production output. Consequently, the principal sets an investment threshold that is progressively lower than the first-best, resulting in overinvestment at a time-varying degree.

The simple form of the optimal investment policies x^{FB} and $x(W)$ summarized in Lemma 1 and Proposition 3 allows the derivation of useful comparative statics, offering empirically relevant predictions testable in the context of both internal and external investments and project selections. According to Corollary 1, agency frictions in the model lead to overinvestment, the degree of which can be measured by the ratio between $x(\underline{W})$ and x^{FB} (i.e., $\theta_{\min}/x^{FB}$). A parameter change is said to exacerbate overinvestment if it lowers this ratio. Based on these definitions, Lemma 1 and Proposition 3 imply the following:

Proposition 4 *A higher intensity of target arrival, λ, or a lower search cost, δ, exacerbate the degree of overinvestment as measured by the ratio between $x(\underline{W})$ and x^{FB}.*

Empirically, the intensity of target arrival λ can be proxied by the number of firms in the industry and/or the frequency of SPACs, M&A, VC, HFA, or internal innovation. The search cost δ can be proxied by the type of industry (more or less innovative), geographical proximity/location (Glaeser and Lang, 2023), executive connections, as well as by standard measures for the availability of information about investment targets (e.g., percentage of public firms in the industry, stock liquidity, analyst coverage, and institutional holdings), given that more information facilitates searching. The degree of overinvestment can be approximated by market reactions—such as stock returns—to the investment decisions (e.g., Lang, Stulz, and Walkling, 1991; Harford, 1999; Andrade, Mitchell, and Stafford, 2001; Masulis, Wang, and Xie, 2007, 2009), accounting measures needed to maintain assets in place and finance positive NPV projects (e.g., Richardson, 2006; Rozenbaum, 2019), or deviations from that predicted by the firm's investment opportunity (e.g., McNichols and Stubben, 2008). Together, the results in Proposition 4 can be translated into the following testable empirical prediction:

Prediction 1 *Overinvestment in SPACs, M&A, VC, HFA, or internal innovation is positively associated with the number of firms and frequency of activities in the relevant markets,* *the geographical proximity between the investment firms and their targets, as well as with the executive connections, stock liquidity, analyst coverage, and institutional holdings of the targets.*

Our model also predicts that along the equilibrium path, the principal optimally adopts a progressively lower investment threshold with values between $x(W^*)$ and $x(\underline{W})$. Because a higher investment threshold is associated with a higher expected payoff ($V_2'(x) > 0$), the dynamics of x can be potentially proxied by the variations in returns. In particular, the gap between $x(W^*)$ and $x(\underline{W})$ can be interpreted as the "return dispersion" in those markets, which is arguably straightforward to measure. Thus, Propositions 3 and 4 suggest that such dispersion is wider if λ is higher or if δ is lower, implying the following hypothesis:

Prediction 2 *The return dispersion of SPAC business combinations, M&A deals, VC investments, HFA targets, internal investments in innovation is positively associated with the frequency of those activities and the number of firms in the market. The return dispersion is negatively associated with the geographical proximity between the investment firms and their targets, as well as with the executive connections, stock liquidity, analyst coverage, and institutional holdings of the targets.*

Our results provide a different perspective for interpreting cross-sectional survey data, such as those reported in Poterba and Summers (1995) and Drury and Tayles (1996). These surveys reveal that firms actively and intentionally forgo positive NPV projects—a practice typically regarded as a sign of underinvestment. Our model suggests that this is not necessarily the case if firms' investment decisions are made in a dynamic environment: In our model, every project has a positive NPV, yielding a strictly positive payoff to the principal even after taking into account the necessary information rents to the agent (i.e., $y(\theta) - w(\theta) > 0$ for all θ). However, the firm may optimally forgo some earlier targets in anticipation of better ones in the future. As time progresses, the firm also optimally lowers its quality standard and may eventually invest in a target with lower returns than those of previously foregone targets. Therefore, without controlling for the time-series of investment opportunities, the cross-sectional observations of investment behaviors alone do not necessarily reflect suboptimal capital budgeting policies.

We further complement the analytical comparative statics with numerical simulations that yield predictions regarding the distribution of the variables of interest. Specifically, for each set of parameters, we simulate 1,000 paths of evolutions of the contract and calculate the success rate (i.e., the fraction of paths in which a suitable target arrives and triggers the investment) and the agent's initial time limit (i.e., the maximum search time allowed before termination). We also calculate the average and standard deviations of the search time, target value, and managerial compensation conditional on the investment being triggered.

Table 1 presents the results for a benchmark case and several comparative statics in which we maintain the value for all but one parameter of the benchmark. A higher λ (target arrival rate) increases both the frequency of deal completion and the maximal search time allowed before termination. In contrast, a higher δ (search cost) or a higher κ, which means a thinner tail for the underlying distribution of target quality (i.e., fewer high-quality targets), both lower the completion rate and the maximal time allowed.

Table 1: **Simulation**

	(1) Benchmark	(2) Higher λ	(3) Higher δ	(4) Higher κ
Success rate	0.36	0.39	0.29	0.28
Initial time limit	272	375	119	98
Conditional on search being successful				
Average search time	173.77	237.61	73.40	59.83
SD of search time	80.35	112.73	33.34	27.31
Average target value	3.71	4.32	3.04	2.68
SD of target value	3.54	3.75	3.54	3.09
Average agent compensation	4.10	5.11	3.37	2.66
Average cost of overinvestment	12.76	14.97	8.86	5.36

The parameters for the benchmark are $\lambda = 3$, $\delta = 1.1$, $\kappa = 2.5$. In columns (2) to (5), all parameters are the same as those in the benchmark except for: $\lambda = 3.25$ in column (2); $\delta = 1.3$ in column (3); and $\kappa = 2.6$ in column (4). Each column corresponds to 1,000 paths of simulations. Success rate is the fraction of the paths in which a suitable target according to the optimal investment policy arrives and investment is triggered. The initial time budget is the maximal search time allowed before contract termination. Managerial compensation of each deal refers to $W_{\tau^-} + R(\theta)$, i.e., the residual utility carried over from the search stage plus the agent's rent in the production stage based on the quality of the target. The average and standard deviations of the search time, target value, average agent compensation and cost of investment are conditional moments of all paths that do not end with termination.

Higher λ implies a *longer* average search time conditional on the search being successful.

There are two reasons for this outcome. First, higher λ increases the maximal allowed search time. Second, with more abundant potential targets, the optimal contract imposes a higher initial hurdle for investment that also declines slowly over time. This can be seen in the higher average target value and managerial compensation, which are determined by the investment hurdle x. In comparison, the agent is given a shorter time budget as well as a more rapidly declining investment hurdle when δ or κ are higher, resulting in a lower average target value and managerial compensation but also a faster search time. These results can be summarized in the following testable prediction, which may help reconcile the empirically observed correlation between the success rate and performance of various search processes, such as the large numbers of SPAC business combinations completed in recent years and their poor subsequent returns (e.g., Gahng, Ritter, and Zhang, 2022).

Prediction 3 *The average returns to investments in SPAC business combinations, M&A deals, VC investments, HFA targets, internal investments in innovation are positively correlated with geographical proximity, frequency of deal completion, number of firms and relative frequency of public firms, executive connections, stock liquidity, analyst coverage, and institutional holdings of the targets, and negatively correlated with the average incentive power of contracts.*

Finally, the simulation reveals a relationship between the cost of overinvestment and managerial compensation. Here, we define

$$\mu^{FB} = \mathrm{E}[\theta|\theta > x^{FB}] = \frac{\kappa x^{FB}}{\kappa - 1} \tag{28}$$

as the expected quality of the targets being invested under the first-best scenario. For each target θ_j chosen in simulation j, we define the "cost of overinvestment" as $\mu^{FB} - \theta_j$ if $\theta_j < x^{FB}$, and zero otherwise. This cost reflects that, due to the agency friction, an inferior target θ_j is chosen instead of waiting for a target that clears the first-best investment hurdle, which is μ^{FB} in expectation. We then take the average of this cost across all paths under which an investment is made in Table 1, which shows that

Prediction 4 *The expected cost of overinvestment is positively correlated with the expected managerial compensation.*

Intuitively, while the manager's compensation increases with the average quality of the targets invested, the cost of overinvestment is driven by the gap between the average quality of the targets invested and the average quality of the targets that should be invested in under the first-best scenario. Our numerical simulations reveal that when a change in the parameters raises the average quality of both types of targets, it also widens the gap between them, thereby increasing the cost of overinvestment. Our prediction corroborates with existing empirical evidence from capital markets. For example, using a comprehensive hand-collected sample of SPACs, Feng, Nohel, Tian, Wang, and Wu (2023) shows that the deals in which the sponsors receive higher compensation are on average associated with poorer performance after the completion of the business combination.

Empirical studies (e.g., Richardson, 2006; Masulis, Wang, and Xie, 2007, 2009; Franzoni, 2009) usually attribute the evidence of overinvestment to the manager's empire building preference and interpret it as a lack of discipline inside the firm. That is, overinvestment in these studies is a *deviation from optimality*, and a higher degree of overinvestment is an indication of poor internal governance. In contrast, we show that overinvestment can arise as an *optimal arrangement* between the firm and the manager even though the manager does not intrinsically derive any additional utility from investment beyond his own compensation. In other words, overinvestment is a necessary feature of optimality. Admittedly, contract design may not always be optimal in practice, and poor internal governance can exacerbate overinvestment. Thus, both mechanisms are likely to coexist in the data. Nevertheless, one might be able to disentangle the two mechanisms by focusing on the group of firms with better internal governance and the ability to adopt optimal contracts.[21] Internal governance can be measured by shareholder protections (e.g., Gompers, Ishii, and Metrick, 2003), board monitoring (e.g., Fich and Shivdasani, 2012; Coles, Daniel, and Naveen, 2014), managerial entrenchment (Bebchuk, Cohen, and Ferrell, 2009). Thus, we can develop the following:

Prediction 5 *Overinvestment is more likely to be the result of optimal contract arrangement*

[21] In the models of Gregor and Michaeli (2020, 2022) eliminating the empire-building bias is too costly so the principal optimally allows some level of overinvestment in equilibrium. There, the party with empire-building preferences is assumed to be in control of the investment policy either directly or indirectly by influencing the information of the decision-maker. Therefore, it may also be possible to determine whether an observed overinvestment is driven by empire-building preferences or not by collecting data on the decision rights inside companies.

if it is observed among firms with better internal governance, stronger board monitoring, and lower managerial entrenchment.

Our predictions are formulated *ceteris paribus*. Empirical testing thus requires identification to control for confounding factors. Although rigorous empirical analysis is outside the scope of this paper, several identification strategies, such as using the decimalization on major stock exchanges as an exogenous shock to stock liquidity (e.g., Edmans, Fang, and Zur, 2013), or the addition of non-stop flights between the locations of a firm and its potential investment targets as an exogenous shock to search cost (e.g., Bernstein, Giroud, and Townsend, 2016), already exist in the literature and may provide useful settings to explore the predictive power of the model in this paper.

6 Extensions and Additional Empirical Predictions

This section introduces two extensions: one in which the search stage is subject to the agent's hidden effort (Section 6.1), and one in which the production stage is also dynamic (Section 6.2). These extensions illustrate the robustness of the main results as well as the theoretical flexibility of our baseline model while producing new empirically testable predictions. For ease of exposition, we focus our discussions on the differences in the assumptions and implications compared to the baseline model, and relegate most of the technical analyses of the solution to the Appendix.

6.1 Moral Hazard During the Search Period

In the baseline model, targets arrive randomly via a Poisson process as long as the principal pays the flow search cost δ. In practice, the search for valuable investment opportunities often requires costly and unobservable effort from the managers, thus generating a moral hazard problem that can interact with the adverse selection problem regarding the timing and quality of the arriving targets. In this subsection we extend the baseline model to explore the impact of such moral hazard on the firms' investment policies.

Specifically, we assume that during the search stage, the agent has unobservable control over the utilization of the resources: he can spend the resources on search, in which case the targets arrive via a Poisson jump process N_t with intensity λ as in the baseline model, or alternatively divert the resources to generate private benefit ρ for himself, in which case no target arrives. Note that this is equivalent to having an unobservable binary action $a_t \in \{0, 1\}$ where the agent either works (when spending the resources on search, $a_t = 1$) or shirks (when he diverts them for private benefits, $a_t = 0$). The parameter ρ represents the perks and benefits from actions that are enjoyable to the agent personally but do not contribute to the discovery of investment targets, such as excessive traveling, spending the firm/fund's resources to build a personal reputation or network, and hiring (unqualified) friends and family members. We assume $\rho < \delta$, so shirking is socially inefficient. The second friction is that, as in the main part of the paper, only the agent observes the arrival of targets and their quality θ.

Similar to the baseline model, the contract can be characterized using the agent's continuation utility W_t, which drifts down in the absence of any suitable target. To provide the incentive for search effort, the downward drift of continuation utility must be kept at a rate commensurate to his private benefit if he chooses to shirk. That is, the agent exerts the search effort if and only if

$$J(x_t) \geq \rho, \tag{29}$$

where $J(x_t)$ is still given by (16). Intuitively, the agent faces a tradeoff between working and shirking: the latter yields flow benefit ρdt. However, because no target arrives while he shirks, his continuation utility drifts down at the rate of $J(x_t)dt$. The agent prefers not to shirk if the above-mentioned cost exceeds the benefit, which is captured by inequality (29).

The principal's optimal choice of investment threshold x_t now solves the same HJB equation (19) subject to the IC condition (29). When (29) is slack, the solution follows the same structure as that described in Section 4.2. However, with unobservable search effort, the choice of x_t is constrained. In particular, the left-hand side of the IC condition (29) is decreasing in x, because the arrival rate of high-quality targets decreases faster than the agent's expected rent from those targets. Therefore, there may exist $\overline{W}$ such that the IC condition is binding when $W \geq \overline{W}$. In that region, $x(W)$ is given by the IC condition, or

$$x(W) = \bar{x} \equiv \left[\frac{\lambda\gamma^2\theta_{\min}^{\kappa}}{\rho(\kappa-2)}\right]^{\frac{1}{\kappa-2}}, \quad \text{if } W \geq \overline{W}. \tag{30}$$

In other words, the optimal investment threshold x is constant for a sufficiently high level of W. Setting $x = \bar{x}$ in (21) implies that $\overline{W}$ solves

$$V_1'(\overline{W}) = \frac{\kappa}{2}\left(\frac{1}{\gamma} - \frac{\delta}{\rho}\right). \tag{31}$$

Altogether, the optimal contract during the search stage can be summarized in the following proposition:

Proposition 5 *Under the optimal contract, if the agent's search effort is unobservable, the principal's value function $V_1(W)$ solves the HJB equation (18) subject to the IC condition (29) and the boundary condition $V_1(0) = 0$. If $\gamma\delta < \rho < \lambda\gamma^2\theta_{\min}^2/(\kappa-2)$, then there exist $\{\underline{W}, \overline{W}\}$ that solve (22) and (31), respectively, such that the optimal investment policy $x(W)$ is given by*

$$x(W) = \begin{cases} \theta_{\min}, & \text{if } W < \underline{W} \\ \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right) V_1'(W)\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}, & \text{if } \underline{W} \leq W < \overline{W} \\ \left[\frac{\lambda\gamma^2\theta_{\min}^{\kappa}}{\rho(\kappa-2)}\right]^{\frac{1}{\kappa-2}}, & \text{if } W \geq \overline{W} \end{cases} \tag{32}$$

where $x'(W) > 0$ for all $W \in (\underline{W}, \overline{W})$.

Figure 2 illustrates the value function $V_1(W)$ and the three regions of the optimal investment policy $x(W)$. Compared to the baseline model, unobservable search effort represents an additional agency friction, which intuitively lowers the principal's value function. Meanwhile, $x(W)$ is lower when W is near W^*, implying that *unobservable search effort exacerbates overinvestment.* Intuitively, unobservable search effort imposes an additional constraint (equation 29) that expedites the decline of W in the equilibrium, raising the cost of maintaining a high investment threshold.

The introduction of unobservable search effort brings new insights to our model. Theoret-

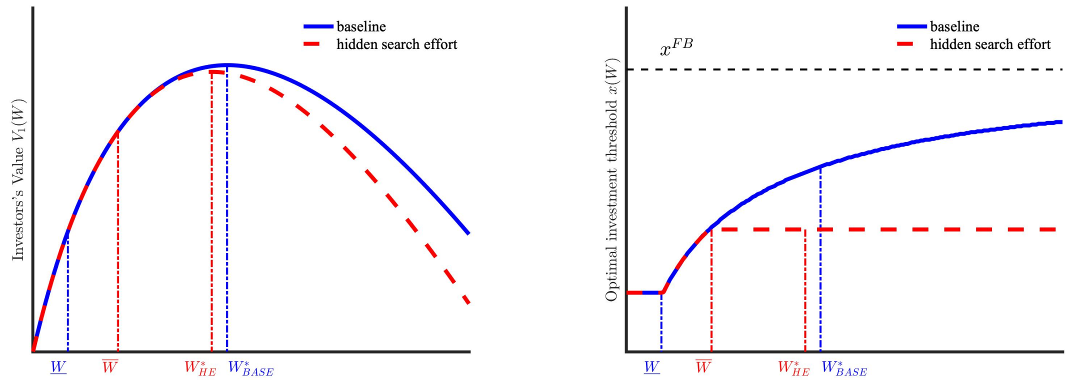

Figure 2: **Value Function and Investment Policy with Unobservable Search Effort** This figure illustrates the incremental effect of unobservable search effort from the agent (red dashed line) compared to the baseline model without such effort (blue solid line). The left panel of this figure plots the principal's value function $V_1(W)$. The right panel plots the optimal investment threshold $x(W)$ and the first-best investment threshold x^{FB}. $\underline{W}$ and $\overline{W}$ are defined according to (22) and (31), respectively, and $W^* \equiv \arg\max_W V_1(W)$ represents the point at which $V_1(W)$ is maximized (W^*_{BASE} for the baseline model and W^*_{HE} for the extension with hidden effort). Parameter values are $\lambda = 2.5$, $\kappa = 4.25$, $\delta = 1.1$, $\rho = 1$.

ically, this extension illustrates a framework in which two widely observed agency frictions—moral hazard and adverse selection—can be jointly analyzed. Azarmsa, Liu, and Noh (2023) studies the incentives for division managers to acquire information at a privately observed cost to influence the allocation policy of the headquarters. There, the main challenge for the headquarters is to maintain the incentive for information acquisition of division managers who are tempted to free-ride the other divisions' information. However, division managers have no incentive to hide or misreport their private information once it is acquired. Our extention differs in that in addition to not exerting the search effort, the agent is also tempted to conceal or misreport his private information (consisting of the arrival and quality of the investment targets). The optimal contract thus must provide incentives not only for the agent's private action but also the truthful and timely report of private information. This connects our study to the broad literature of mechanism design in which the agent can take private actions, such as Kirby, Reichelstein, Sen, and Paik (1991), Dutta and Reichelstein (2002), Krähmer and Strausz (2011), Halac, Kartik, and Liu (2016), and Liu and Lu (2018). In particular, Halac, Kartik, and Liu (2016) studies experiments in a learning model in which the agent's private effort is a necessary (but not sufficient) condition for success. The adverse selection of the underlying success likelihood results in a screening contract with different endogenous deadlines at which the contract is terminated if success has not been achieved. Our extention shares the similarities that "success" (the arrival of an investment target) is stochastic and only possible if the agent exerts effort. However, our model differs in that "success" carries the additional information about the quality and thus requires different levels of incentives under the optimal screening contract. This additional dimension of the agent's private information also implies that the definition of "success" in our model is time-varying: earlier targets must clear a higher hurdle in order to successfully trigger investment. Moreover, while the interaction of the two problems can impose substantial analytical challenges in a general model, our setting allows us to tackle the problems sequentially and thus achieving tractable analytical solutions.

Empirically, the introduction of unobservable search effort produces additional testable hypothesis. In particular, it introduces an additional parameter ρ which measures the agent's private benefit/perk and is related to the principal's optimal investment policies. In the equilibrium, ρ equals the minimal speed at which the agent's continuation utility has to drift down without the target. Therefore, ρ can be indirectly measured by the incentive power of the managerial contract such as the fraction of inside equity. This is also the standard interpretation in the optimal contracting literature (e.g., DeMarzo and Sannikov, 2006; Biais, Mariotti, Plantin, and Rochet, 2007). Thus, we obtain the following hypothesis:

Prediction 6 *The average incentive power of managerial contracts in the investment firms is positively associated with overinvestment in SPACs, M&A, VC, HFA, or internal project selections and innovations.*

6.2 Dynamic Adverse Selection

Our second extension pertains to the production stage. The baseline model posits that production is a one-time decision. Once investment is made and the agent exerts the production effort e, a single output y is realized, and the contracting relationship ends. This implies a static adverse selection problem and simplifies the derivation of the optimal screening contract. In this section, we demonstrate that our main results hold qualitatively when the production stage is also dynamic and the agent's private information evolves stochastically.

Let τ represent the end of the search stage. Consider the following extension: The production stage lasts an exogenous period of $T > 0$ (i.e., from τ to $\tau + T$), during which the agent continuously produces outputs from the target chosen in the previous stage. The production technology is given by

$$y_t = e_t \xi_t, \tag{33}$$

where y_t is the output, e_t is the agent's (production) effort exerted at a quadratic personal cost $h(e_t) = e_t^2/2$, and ξ_t is the "productivity" of the target, which now evolves over time. For tractability, we assume ξ_t follows a geometric Brownian motion (GBM):

$$d\xi_t = \xi_t \left(\mu dt + \sigma dZ_t \right) \tag{34}$$

with publicly-known parameters μ and σ. θ determines the initial value of ξ_t, i.e., $\theta = \xi_\tau$.

Compared to the existing literature, a theoretical innovation (and challenge) of this setting is that the agent's private information ξ_t is *persistent*, which implies the adverse selection problem faced by the principal is *dynamic*. While there are studies exploring persistent private information in the context of dynamic moral hazard (e.g., Williams, 2011; Marinovic and Varas, 2019), studies of persistent private information in the context of adverse selection are rare, as it is known to be a challenging problem. Fortunately, the specific structure of our model implies that the screening problem during the production stage is also *time-separable*. That is, the set of feasible contract terms (i.e., $\{y_t, w_t\}$) at time $t \in [\tau, \tau + T]$ is independent of the history of the contract, and the flow utility of the agent and the principal at time t depends only on the initial and the current private information of the agent. Under time-separability, the dynamic adverse selection problem can be converted into a static mechanism design problem similar to that analyzed in Section 4.1, allowing us to uniquely pin down $R(\theta)$ and $V_2(x)$ under any incentive compatible contract which is all we need to feed back into the search stage. The resulting optimal contract under is summarized as follows:

Proposition 6 *For any given investment policy x, the optimal contract in the production* *stage has the following properties:*

– *The agent's information rent from a target of quality $\theta > x$ at the beginning of the production stage is given by*

$$R(\theta) = \frac{\phi\gamma^2}{2}(\theta^2 - x^2). \tag{35}$$

where ϕ is a function of the model parameters μ, σ, and T, given by (88).

– *The principal's expected payoff at the beginning of the production stage is*

$$V_2(x) = \frac{\phi\gamma\kappa}{2(\kappa-2)}x^2. \tag{36}$$

– *During the production stage, the principal's optimal output target $\{y_t^*\}_{t\in[0,T]}$ is given by $y_t^* = \gamma\xi_t^2$ and the implied equilibrium production effort is $e_t^* = \gamma\xi_t$.*

The optimal contract during the search stage can be summarized by the principal's value function $V_1(W)$, which solves an HJB equation analogous to (18) subject to the boundary condition $V_2(0) = 0$. If $\phi \geq 1$, the optimal investment policy $x(W)$ is given by

$$x(W) = \begin{cases} \theta_{\min}, & \text{if } W \leq \underline{W} \\ \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right) V_1'(W)\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}, & \text{if } W > \underline{W} \end{cases} \tag{37}$$

where $\underline{W}$ is given by (22). $x(W) < x^{FB}$, and $x'(W) > 0$ for all $W > \underline{W}$.

Despite the dynamic nature of the adverse selection problem, our main results remain qualitatively intact. In particular, the agent's information rent in the production stage is still a quadratic function of the target quality θ and the principal's expected payoff is still a quadratic function of the investment threshold x. In the search stage, W drifts downward in the absence of a suitable target, and the optimal investment threshold x is progressively lower, which leads to overinvestment.

The empirical implications of our model variation with dynamic adverse selection are similar to those in the baseline model with the addition of a new one pertaining to the parameter ϕ, which can be interpreted as the industry or regional average return of internal investments and innovations, M&A, HFA, VC, or SPAC activities. It is straightforward to see that a higher ϕ increases x^{FB} while $x(\underline{W}) = \theta_{\min}$ is unchanged, leading to the following testable prediction:

Prediction 7 *The overinvestment in SPACs, M&A, VC, HFA, or internal project selections and innovations is positively associated with the average returns in those markets.*

As noted, the results in this section demonstrate the robustness of the main model and its practical implications. Nevertheless, the solution technique of this dynamic version of the adverse selection problem is far more involved than the one used in the static version and is potentially applicable to a broad class of questions involving persistent and time-varying private information. Our solution method utilizes the Myersonian approach developed in Eső and Szentes (2007) and Pavan, Segal, and Toikka (2014) but extended to continuous time. The details are provided in the proof of Proposition 6 in the Appendix for interested readers. This method can be easily combined with dynamic moral hazard as introduced in Section 6.1, providing a unified framework for researchers interested in jointly studying these two important agency frictions.

7 Concluding Remarks

We consider a setting where a principal delegates a costly and dynamic search for investment targets as well as their operation to an agent. The agent is privy to the (stochastic) arrival of the targets and their quality, and needs to receive proper incentives to truthfully disclose such information. The optimal contract involves the use of a dynamic budget, a threshold for investment, and an internal charge for search resources. The balance of the budget, the investment threshold, and the resource charge are all endogenous and time-varying, adjusted according to the search history. These features yield implications different from those in static models, particularly overinvestment relative to the first-best at an accelerating rate. The results help understand and predict the returns from investments in SPAC business combinations, mergers and acquisitions, VC investments, HFA interventions, and/or internal innovations.

Our work can be extended in several directions. For simplicity, the model does not allow the agent to revert to the search stage once production begins. In practice, the search and production processes do not always move forward linearly. Letting the agent conduct both search and production repeatedly or simultaneously may yield interesting insights about firms' optimal internal organization and/or resource allocation. The agent may be allowed to exert variable levels of effort in order to expedite the search process. Finally, the principal may also have access to an auditing technology that can reveal the quality of the announced target at a cost. We leave these topics for future work.

Appendix

Proof of Lemma 1: First, $V_2^{FB\prime}(\theta) = \theta > 0$. Therefore, if $x \in \Theta^{FB}$ for some x, then $\theta \in \Theta^{FB}$ for all $\theta > x$. i.e., $\Theta^{FB} = \{\theta : \theta \geq x^{FB}\}$. Thus, the conditional expectation of the principal's payoff from setting an arbitrary cutoff investment quality x is

$$\mathcal{V}_2^{FB}(x) \equiv \mathrm{E}\left[V_2^{FB}(\theta)|\theta \geq x\right] = \int_x^{+\infty} \frac{\theta^2}{2}\left(\frac{\kappa x^\kappa}{\theta^{\kappa+1}}\right) d\theta = \frac{\kappa x^2}{2(\kappa-2)}. \tag{38}$$

Equation (38) utilizes the fact that a Pareto distribution truncated from below at some $x > \theta_{\min}$ is a Pareto distribution with the same shape parameter κ and scale parameter x. Equation (38) also reveals why $\kappa > 2$ is needed for the first-best to exist. Let $V_1^{FB}(x)$ be the principal's value function at the outset of the search stage associated with cutoff policy x, then

$$V_1^{FB}(x) = \max_x \int_0^{+\infty} \left[-\delta t + F(x)V_1^{FB}(x) + (1-F(x))\mathcal{V}_2^{FB}(x)\right] \lambda e^{-\lambda t} dt. \tag{39}$$

The three terms inside the square brackets represent the cost of search, the payoff from a target with quality lower than x, and the expected payoff from the arrival of a target with quality x or above, respectively. Using the fact that $F(x) = 1 - (\theta_{\min}/x)^\kappa$ for Pareto distribution, $V_1^{FB}(x)$ satisfies the following Hamilton-Jacobi-Bellman (HJB) equation:

$$0 = \max_x \ -\delta + \lambda\left(\frac{\theta_{\min}}{x}\right)^\kappa \left[\mathcal{V}_2^{FB}(x) - V_1^{FB}(x)\right]. \tag{40}$$

Substituting (38) into the HJB equation and re-arranging the terms yields

$$V_1^{FB}(x) = \max_x \ \frac{\kappa x^2}{2(\kappa-2)} - \frac{\delta x^\kappa}{\lambda \theta_{\min}^k}. \tag{41}$$

The first order condition with respect to x yields x^{FB} as in (4). Substituting x^{FB} into (41) yields $V_1^{FB}(x^{FB})$ as in (5).

Proof of Lemma 2: Suppose there is an incentive-compatible optimal contract $\mathcal{C}$ under which an open set H exists with the following properties: for all $\theta \in H$, $\theta \notin \Theta$, and there exists θ' such that $\theta' \in \Theta$ but $\theta' < \tilde{\theta} \equiv \inf H$. Let $\bar{\theta} \equiv \max\{\theta : \theta \in \Theta, \theta < \tilde{\theta}\}$. Clearly, $y(\bar{\theta}) - w(\bar{\theta}) > 0$ and $w(\bar{\theta}) \geq h(e(\bar{\theta}))$ if $\mathcal{C}$ is optimal. Now consider a contract $\mathcal{C}'$ that is otherwise identical to $\mathcal{C}$ except for the following: for any report $\widehat{\theta} \in H$, the required output $y(\widehat{\theta}) = y(\bar{\theta})$ and the associated compensation is $w(\widehat{\theta}) = w(\bar{\theta})$. This contract is incentive-compatible because for all $\tilde{\theta} \in H$, $e(\tilde{\theta}) = y(\bar{\theta})/\tilde{\theta} < e(\bar{\theta})$ and is independent of the report $\widehat{\theta}$. Thus, $w(\tilde{\theta}) = w(\bar{\theta}) > h(e(\tilde{\theta}))$ for all $\tilde{\theta}$. However, $y(\bar{\theta}) - w(\bar{\theta}) > 0$ implies that $\mathcal{C}'$ generates the same payoff as $\mathcal{C}$ for all $\theta \notin H$ but positive (higher) payoff for all $\tilde{\theta} \in H$, contradicting the assumption that $\mathcal{C}$ is optimal. Therefore, it must be that such H does not exist under the optimal contract.

Proof of Proposition 1: The principal offer a screening contract $\{w(\widehat{\theta}), y(\widehat{\theta})\}$. The agent reports his type $\widehat{\theta}$, produces the required level of output, and receives the associated compensation. With a slight abuse of notation, define

$$R(\theta, \widehat{\theta}) = w(\widehat{\theta}) - h(e) \tag{42}$$

as the information rent of the agent with type-θ reporting $\widehat{\theta}$, subject to the constraint that $y(\theta, e) = y(\widehat{\theta})$. i.e., he must produce the level of output designed for the type-$\widehat{\theta}$ agent. Let $e(y, \theta)$ represent the necessary effort required by a type-θ agent to produce output y. Then, one can define $R(\theta) = R(\theta, \theta)$ as the agent's equilibrium rent under truthful reporting, and

$$\widehat{\theta}^*(\theta) = \arg\max_{\widehat{\theta}} \; R(\theta, \widehat{\theta}) \tag{43}$$

as the optimal reported type chosen by a type-θ agent. This optimality implies the following envelope condition:

$$R_{\widehat{\theta}}(\theta, \widehat{\theta}^*(\theta)) = 0. \tag{44}$$

Therefore, in the equilibrium

$$R'(\theta) = \frac{\partial R(\theta, \widehat{\theta}^*(\theta))}{\partial \theta} = R_\theta + R_{\widehat{\theta}}(\theta, \widehat{\theta}^*(\theta)) \frac{d\widehat{\theta}^*(\theta)}{d\theta} = R_\theta = -h'(e) e_\theta(y, \theta) \tag{45}$$

based on the envelope condition.

The principal's payoff in the production stage is therefore $V_2(x) - W_{t^-}$, where

$$V_2(x) = \max_{y,w} \int_x^\infty [y(\theta) - w(\theta)] dF(\theta) = \max_{y,w} \int_x^\infty [y(\theta) - h(e(\theta, y)) - R(\theta)] dF(\theta) \tag{46}$$

where

$$F(\theta) = 1 - \left(\frac{x}{\theta}\right)^\kappa \tag{47}$$

$$f(\theta) = \frac{\kappa x^\kappa}{\theta^{\kappa+1}}. \tag{48}$$

Applying integration by parts to the last term inside the integral of $V_2(x)$ in (46) yields

$$\int_x^\infty R(\theta) dF(\theta) = \int_x^\infty R'(\theta)(1 - F(\theta)) d\theta + R(x). \tag{49}$$

Substituting this into (46) above yields

$$V_2(x) = \max_y \int_x^\infty [y - h(e) - R'(\theta) g(\theta)] f(\theta) d\theta \tag{50}$$

where

$$g(\theta) \equiv \frac{1 - F(\theta)}{f(\theta)}$$

represents the inverse hazard function of θ. Replacing $R'(\theta)$ with (45), point-wise maximization with respect to y yields the following optimality condition:

$$1 - h'(e)e_y(e,\theta) + \frac{dh'(e)e_\theta(y,\theta)}{dy}g(\theta) = 0 \tag{51}$$

which yields the optimal target y for each type θ. Because $y = \theta e$ and $h(e) = e^2/2$, $e(y,\theta) = y/\theta$, $e_\theta = -y/\theta^2$, and

$$\frac{dh'(e)e_\theta(y,\theta)}{dy} = -\frac{dy^2/\theta^3}{dy} = -\frac{2y}{\theta^3}.$$

The fact that θ follows a Pareto distribution implies that

$$g(\theta) = \frac{1 - F(\theta)}{f(\theta)} = \frac{\theta}{\kappa} \tag{52}$$

Substituting these results into (51) yields

$$1 - \frac{y}{\theta^2} - \frac{2y}{\kappa\theta^2} = 0 \tag{53}$$

which implies $y = \gamma\theta^2$, where

$$\gamma = \frac{\kappa}{\kappa + 2}. \tag{54}$$

Substituting $y = \gamma\theta^2$ into (42), using the IC constraint $\widehat{\theta} = \theta$, and imposing $R(x) = 0$ yields

$$R(\theta) = \frac{\gamma^2}{2}(\theta^2 - x^2). \tag{55}$$

Combine this with $y = \gamma\theta^2$ implies that

$$V_2(x) = \frac{\gamma\kappa}{2(\kappa - 2)}x^2. \tag{56}$$

Proof of Proposition 2: Let $\mathcal{F}_t$ denote the filtration generated by the agent's report θ_t (where $\theta_t = 0$ if no investment opportunity arrives). W_t is an $\mathcal{F}_t$-martingale and thus, by the martingale representation theorem for jump processes, there exists a $\mathcal{F}_t$-predictable, integrable process β_t such that

$$dW_t = a_t\beta_t(dN_t - \lambda(1 - F(x_t))dt). \tag{57}$$

Incentive compatibility of truthful reporting of θ requires that $W_\tau - W_{\tau^-} = R(\theta_\tau)$ if the contract moves to the next stage, which implies $\beta_t = \mathrm{E}[R(\theta_t)|\theta_t \geq x_t] = U(x_t)$ by the property of a martingale.[22] Thus, given the investment policy x_t,

$$dW_t = U(x_t)(dN_t - \lambda(1 - F(x_t))dt) \tag{58}$$

where $U(x_t)\lambda(1-F(x_t)) \geq \rho$. Substituting in $1 - F(x_t) = \left(\frac{\theta_{\min}}{x_t}\right)^{\kappa}$ yields equations (16) and (29).

Proof of Proposition 3: Applying Ito's lemma to dW_t implies the principal's value function in the search stage solves the following HJB equation:

$$0 = \max_{x} \quad -\delta - \lambda \left(\frac{\theta_{\min}}{x}\right)^{\kappa} U(x)V_1'(W) + \lambda \left(\frac{\theta_{\min}}{x}\right)^{\kappa} [V_2(x) - W - V_1(W)] \tag{59}$$

Substituting $U(x)$ from (13) and $V_2(x)$ from (12) into the HJB equation and rearranging terms yields:

$$V_1(W) = \max_{x} \quad \left(\frac{\gamma\kappa}{2(\kappa-2)}\right) x^2 - \left(\frac{\gamma^2 x^2}{\kappa - 2}\right) V_1'(W) - W - \frac{\delta}{\lambda} x^{\kappa} \theta_{\min}^{-\kappa}. \tag{60}$$

The first order condition implies

$$\frac{\gamma}{\kappa-2}\left[\kappa - 2\gamma V_1'(W)\right] x = \frac{\delta}{\lambda} \kappa x^{\kappa-1} \theta_{\min}^{-\kappa}. \tag{61}$$

The solution is

$$x(W) = \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right) V_1'(W)\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}. \tag{62}$$

Meanwhile, substituting $x = \theta_{\min}$ into (62) implies that there exists $\underline{W}$ such that

$$\theta_{\min} = \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right) V_1'(\underline{W})\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}. \tag{63}$$

That is, $\underline{W}$ solves

$$V_1'(\underline{W}) = \left(1 - \frac{\delta(\kappa-2)}{\gamma\lambda\theta_{\min}^2}\right)\frac{\kappa}{2\gamma}. \tag{64}$$

Note that under the optimal contract, $V'(W) > -1$ for all W. This is because principal can always make a cash transfer to the agent, which lowers W and $V(W)$ by the exact same amount. Therefore, the marginal value of building W inside the firm can never be lower than the marginal value of cash transfer, which is -1.

[22] Note that this also follows the definition of $R(\theta)$ in Section 4.1 which is the information rent the agent must be given to reveal θ truthfully *in addition to* any utility W_{τ^-} carried over from the search stage.

The proof above assumes $R(x) = 0$: i.e., it is optimal not to award the agent additional compensation beyond that necessary for the truthful revealing of target quality. It is straightforward to show that this is indeed optimal. Let $b(W) \equiv R(x(W))$ denote such the extra bonus. The HJB equation (19) can be then re-written as

$$V_1(W) = \max_{x,b} \quad V_2(x) - b - W - (U(x) + b)V_1'(W) - \frac{\delta}{\lambda}\left(\frac{x}{\theta_{\min}}\right)^{\kappa} \tag{65}$$

subject to $b \geq 0$ and $x \geq \theta_{\min}$. The first order derivative for b is $-1 - V_1'(W) < 0$. Consequently, $b = 0$ for all W. Intuitively, a bonus payment decreases V_2 while increasing the drift of dW_t, both of which are costly to the principal.

Proof of Corollary 1: Note that under the optimal contract, $V'(W) > -1$ for all W. This is because principal can always make a cash transfer to the agent, which lowers W and $V(W)$ by the exact same amount. Therefore, the marginal value of building W inside the firm can never be lower than the marginal value of cash transfer, which is -1. Substituting $V'(W) = -1$ into the first-order condition (21) implies that

$$\lim_{W\to+\infty} x(W) = \left[\gamma\left(1 + \frac{2\gamma}{\kappa}\right)\left(\frac{\lambda\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}} \tag{66}$$

Because

$$\frac{2\gamma}{\kappa} = \frac{2\kappa}{\kappa(\kappa+2)} < 1 \tag{67}$$

when $\kappa > 2$, this limit is always smaller than x^{FB}. Thus, $x(W) < x^{FB}$ for all W.

Proof of Proposition 4: Equations (4) and (30) imply the following comparative statics regarding x^{FB} and $x(\overline{W})$.

$$\frac{\partial x^{FB}}{\partial \delta} < 0, \quad \frac{\partial x^{FB}}{\partial \lambda} > 0. \tag{68}$$

Given that $x(\underline{W}) = \theta_{\min}$ is a constant, these results imply that

$$\frac{\partial}{\partial \lambda}\left(\frac{x(\underline{W})}{x^{FB}}\right) < 0, \quad \frac{\partial}{\partial \delta}\left(\frac{x(\underline{W})}{x^{FB}}\right) > 0. \tag{69}$$

Proof of Proposition 5: The proof is similar to that of Proposition 3. Thus, we only highlight the differences. First, due to the hidden search effort, the agent's continuation utility W_t is now defined as

$$W_t = \mathrm{E}\left[\int_t^{\tau} \rho(1-a_s)ds + \int_t^{\tau} dC_s + W_{\tau}\right]. \tag{70}$$

Compared to (15), the new term $\rho(1-a_s)ds$ captures the agent's private benefit if he shirks. It is straightforward to show that the optimal contract always implements no shirking (i.e., $a_t = 1$) during the search stage. This is because search requires a cost δ but shirking generates a benefit $\rho < \delta$ to the agent. Therefore, any contract that involves shirking can be strictly improved by discouraging shirking through compensating the agent for his lost shirking benefit.

The principal solves the HJB equation (19), subject to the IC constraint (29). If this constraint is slack, the solution is identical to that in Proposition 3. If (29) is binding, then

$$\lambda\left(\frac{\theta_{\min}}{x}\right)^{\kappa}U(x) = \frac{\lambda\theta_{\min}^{\kappa}}{x^{\kappa}}\left(\frac{\gamma^2 x^2}{\kappa-2}\right) = \rho \tag{71}$$

yields the solution

$$x(W) = \bar{x} \equiv \left[\frac{\lambda\gamma^2\theta_{\min}^{\kappa}}{\rho(\kappa-2)}\right]^{\frac{1}{\kappa-2}}. \tag{72}$$

Substituting (72) into (62) implies that there exists $\overline{W}$ such that

$$x(\overline{W}) = \left[\left(1-\left(\frac{2\gamma}{\kappa}\right)V_1'(\overline{W})\right)\left(\frac{\lambda\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}} = \bar{x}. \tag{73}$$

That is, $\overline{W}$ solves

$$V_1'(\overline{W}) = \frac{\kappa}{2}\left(\frac{1}{\gamma}-\frac{\delta}{\rho}\right). \tag{74}$$

The existence of both $\{\underline{W}, \overline{W}\}$ requires that $\overline{W} > \underline{W}$, which is equivalent to

$$\frac{\kappa}{2}\left(\frac{1}{\gamma}-\frac{\delta}{\rho}\right) < \left(1-\frac{\delta(\kappa-2)}{\gamma\lambda\theta_{\min}^2}\right)\frac{\kappa}{2\gamma} \tag{75}$$

which simplifies to

$$\lambda\theta_{\min}^2 > \frac{\rho(\kappa-2)}{\gamma^2}; \tag{76}$$

Meanwhile, $V_1'(\overline{W}) > 0$ requires that

$$\frac{\kappa}{2}\left(\frac{1}{\gamma}-\frac{\delta}{\rho}\right) > 0 \tag{77}$$

which implies

$$\rho > \gamma\delta. \tag{78}$$

Finally, we prove that $R(x) = 0$ is still optimal for all $W \leq W^*$ (i.e., along the equilibrium path). The HJB equation (19) can be then re-written as

$$V_1(W) = \max_{x,b} \quad V_2(x) - b - W - (U(x) + b)V_1'(W) - \frac{\delta}{\lambda}\left(\frac{x}{\theta_{\min}}\right)^{\kappa} \tag{79}$$

subject to the modified IC constraint

$$\left[\lambda\left(\frac{\theta_{\min}}{x}\right)^{\kappa}\right][U(x) + b] \geq \rho \tag{80}$$

and that $b \geq 0$ and $x \geq \theta_{\min}$. Let π be the Lagrangian multiplier associated with the modified IC constraint. The first order derivative for b is

$$-1 - V_1'(W) + \pi\lambda\theta_{\min}^{\kappa}x^{-\kappa}. \tag{81}$$

If (80) is slack (i.e., $\pi = 0$), then $b = 0$ is clearly optimal since $V_1'(W) > -1$. If (80) is binding (i.e., $\pi > 0$) for any $W \leq W^*$, it must be that (78) holds. Since $V_1'(W) \geq 0$ and $V_1''(W) < 0$, it is sufficient to show that $b = 0$ is optimal when $W = W^*$, where the first order condition for x implies

$$\frac{\gamma\kappa x}{\kappa - 2} - \frac{\delta}{\lambda}\kappa x^{\kappa-1}\theta_{\min}^{-\kappa} + \pi\lambda\theta_{\min}^{\kappa}\left[\left(\frac{2\gamma}{\kappa - 2}\right)x^{1-\kappa} + bx^{-1-\kappa}\right] = 0. \tag{82}$$

Evaluating this at $b = 0$ and $x = \bar{x}$ where $\bar{x}$ is given by (30) yields

$$\pi\lambda\theta_{\min}^{\kappa}x^{-\kappa} = \left(\frac{\delta}{\lambda}\kappa\bar{x}^{\kappa-2}\theta_{\min}^{-\kappa} - \frac{\gamma\kappa}{\kappa - 2}\right)\left(\frac{\kappa - 2}{2\gamma}\right) = \frac{\kappa}{2}\left(\frac{\delta}{\rho} - 1\right). \tag{83}$$

Thus,

$$-1 + \pi\lambda\theta_{\min}^{\kappa}x^{-\kappa} = -1 + \frac{\kappa}{2}\left(\frac{\delta}{\rho} - 1\right) < 0 \tag{84}$$

since (78) implies $\delta/\rho < 1/\gamma = 1 + 2/\kappa$. Intuitively, the marginal benefit of setting $b > 0$ is to relax the IC constraint (80). However, this is insufficient compared to its marginal costs, including a higher $U(x)$, a lower $V_2(x)$, and a higher search cost paid. Therefore, setting $b = 0$ is optimal along the equilibrium path.

Proof of Proposition 6: This proof is organized into three subsections.

A. Preliminaries

When ξ_t follows a GBM, $\xi_t = \xi_\tau \nu_t$, where

$$\nu_t = \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma Z_t\right], \quad (\nu_\tau = 1) \tag{85}$$

is an exogenous stochastic process with known distribution for any given t. Therefore, $d\xi_t/d\theta = \xi_t/\xi_\tau = \nu_t$. Put differently, the marginal value of target quality θ on its subsequent productivity at any time during the production stage depends on the path of exogenous shocks only, a property that greatly simplifies the analysis below.

If all information is public, the first-best effort and output in the production stage solve

$$\max_{e_t} \; y_t - h(e_t) = e_t \xi_t - h(e_t). \tag{86}$$

The solution is $e_t^{FB} = \xi_t$, $y_t^{FB} = \xi_t^2$. The principal's expected payoff given the quality of the target, $V_2^{FB}(\theta)$, is given by

$$V_2^{FB}(\theta) = \mathrm{E} \int_{\tau}^{\tau+T} \left(y_t^{FB} - \frac{1}{2} \left(e_t^{FB} \right)^2 \right) dt = \frac{\phi}{2} \theta^2, \tag{87}$$

where

$$\phi = \mathrm{E} \int_{\tau}^{\tau+T} \nu_t^2 dt = \frac{e^{(2\mu+\sigma^2)T} - 1}{2\mu + \sigma^2} \tag{88}$$

is a constant. Here, ϕ measures the marginal value of target quality summarizing the joint effect of μ, σ, and T, and will be treated as a known parameter in the subsequent analysis.

When ξ_t (including θ) is the agent's private information, let $\widehat{\theta}$ and $\{\widehat{\xi_t}\}_{t \in (\tau, \tau+T]}$ represent the agent's *reported* target quality and time-t productivity, respectively. The resulting screening contract now involves a series of the output target $\{y_t^{\widehat{\theta}}(\widehat{\xi_t})\}_{t \in (\tau, \tau+T]}$ and the corresponding wage $\{w_t^{\widehat{\theta}}(\widehat{\xi_t})\}_{t \in (\tau, \tau+T]}$ if the required output is produced. Conditional on any utility W_{τ^-} carried over from the search stage, the agent's objective is to maximize his expected wage minus his (production) effort cost from the contract—his information rent—which is given by

$$R(\theta) = \max_{\widehat{\theta}, \widehat{\xi}_t, e_t} \; \mathrm{E} \left[\int_{\tau}^{\tau+T} (w_t^{\widehat{\theta}}(\widehat{\xi_t}) - h(e_t)) dt \right], \tag{89}$$

subject to the constraint that $e_t \xi_t = y_t^{\widehat{\theta}}(\widehat{\xi_t})$. The principal's objective is to maximize her payoff at the outset of the production stage, which is $V_2(x) - W_{\tau^-}$, where

$$V_2(x) = \max_{y_t, w_t} \int_x^{+\infty} \mathrm{E} \left[\int_{\tau}^{\tau+T} (y_t - w_t) dt \right] dF(\theta), \tag{90}$$

subject to the IC constraint $\widehat{\theta} = \theta$ and $\widehat{\xi_t} = \xi_t$ for all $t \in (\tau, \tau + T]$.

B. The Production Stage

This section solves the optimal screening contract in the production stage. The proof begins with deriving the agent's information rent for *any* incentive compatible contract. Consider any report $\widehat{\theta}$ made by an agent possessing an arbitrary θ-quality target. Based on this report, the agent is assigned the contract $\mathcal{C}(\widehat{\theta})$, which imposes output target $y^{\widehat{\theta}}(\widehat{\xi_t})_t$ for any future report $\widehat{\xi_t}$, the associated wage $w^{\widehat{\theta}}(\widehat{\xi_t})_t$ if the required output is produced. The contract implies a recommended effort process $\widehat{e}_t \equiv y_t^{\widehat{\theta}} / \widehat{\xi_t}$ for all t. Therefore, given the true productivity process ξ_t, the agent's actual effort choice e_t must satisfy:

$$e_t \xi_t = \widehat{e}_t \widehat{\xi}_t. \tag{91}$$

The payoff for the agent is

$$R(\theta; \widehat{\theta}) = \mathrm{E}\left[\int_{\tau}^{\tau+T} \left(w^{\widehat{\theta}}(\widehat{\xi}_t) - h(e_t)\right) dt\right]. \tag{92}$$

In principle, $\widehat{\theta}$ represents a very large set of possible deviations by the agent. However, Pavan, Segal, and Toikka (2014) and the subsequent studies of dynamic adverse selection (e.g., Bergemann and Strack, 2015, Gershkov, Moldovanu, and Strack, 2018, etc.) show that if the screening problem is time-separable, it is without the loss of generality to establish the IC condition for a particular type of deviation: if he misreports the target quality, $\widehat{\theta} \neq \theta$, his follow-up strategy is to continue misreporting *as if the true quality was* $\widehat{\theta}$ *and he had reported that truthfully*. More formally, at any time, the agent's reported productivity satisfies the following so-called *consistent deviation*:

$$\widehat{\xi}_t = \widehat{\theta} v_t = \widehat{\theta} \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right) t + \sigma Z_t\right] \tag{93}$$

where Z_t represents the true productivity shocks the misreporting agent experiences. This implies that although the agent's private information regarding ξ_t is persistent, it is without loss of generality to label each agent only by the quality of his target θ. Thus, we can differentiate (92) with respect to θ to obtain:

$$\frac{\partial}{\partial \theta} R(\theta; \widehat{\theta}) = \mathrm{E}\left[\int_{\tau}^{\tau+T} \left(-\frac{\partial}{\partial \theta} h(e_t)\right) dt\right] \tag{94}$$

$$= \mathrm{E}\left[\int_{\tau}^{\tau+T} \left(\frac{e_t \widehat{e}_t \widehat{\xi}_t}{\theta \xi_t}\right) dt\right] \tag{95}$$

$$= \mathrm{E}\left[\int_{\tau}^{\tau+T} \left(\frac{e_t \widehat{e}_t \widehat{\theta}}{\theta^2}\right) dt\right] \tag{96}$$

where the second line utilizes the constraint (91), and the third line utilizes the consistent deviation (93). Evaluating (96) at the equilibrium ($\widehat{e}_t = e_t, \widehat{\xi}_t = \xi_t$) and substituting e_t with y_t/ξ_t implies

$$R'(\theta) = \mathrm{E}\left[\int_{\tau}^{\tau+T} \left(\frac{e_t^2}{\theta}\right) dt\right] = \mathrm{E}\left[\int_{\tau}^{\tau+T} \frac{1}{\theta} \left(\frac{y_t}{\xi_t}\right)^2 dt\right] \tag{97}$$

which is the *dynamic envelope condition* analogous to the envelope condition derived in the proof of Proposition 1 above. Integrating (97) from x up yields the information rent for any given θ-quality target:

$$R(\theta) = \int_x^\theta \mathrm{E}\left[\int_\tau^{\tau+T} \frac{1}{q}\left(\frac{y_t}{\xi_t}\right)^2 dt\right] dq + R(x). \tag{98}$$

Next, we derive the principal's expected payoff given the distribution of θ. With a slight abuse of notation, let $\int_x^\infty (\cdot) dF(\theta; x)$ denote the expectation of θ taken under the support Θ taking into account how the distribution of $F(\theta)$ shifts with x. The principal's maximal payoff at the outset of the production stage under any incentive compatible contract is $V_2(x) - W_{\tau^-}$, where

$$V_2(x) = \max_{y_t, w_t} \int_x^\infty \mathrm{E}\left[\int_\tau^{\tau+T} (y_t - w_t) dt\right] dF(\theta; x). \tag{99}$$

The definition of information rent R (Eq. 89) implies that

$$\mathrm{E}\left[\int_\tau^{\tau+T} w_t dt\right] = R(\theta) + \mathrm{E}\left[\int_\tau^{\tau+T} h(e_t) dt\right]. \tag{100}$$

Substituting this into the definition of $V_2(x)$ (Eq. 99) yields

$$V_2(x) = \max_{y_t} \int_x^\infty \mathrm{E}\left[\int_\tau^{\tau+T} \left(y_t - \frac{1}{2}\left(\frac{y_t}{\xi_t}\right)^2\right) dt\right] dF(\theta; x) - \int_x^\infty R(\theta) dF(\theta; x). \tag{101}$$

Applying integration by parts and the fundamental theorem of calculus to the last term yields

$$\int_x^\infty R(\theta) dF(\theta; x) = \int_x^\infty R'(\theta)\left(\frac{1 - F(\theta; x)}{f(\theta; x)}\right) dF(\theta; x) + R(x) \tag{102}$$

$$= \int_x^\infty R'(\theta)\left(\frac{\theta}{\kappa}\right) dF(\theta; x) + R(x) \tag{103}$$

where the second line comes from the property of the Pareto distribution. Clearly, $R(x) = 0$ under the optimal contract. Replacing $R'(\theta)$ with (97) and substituting the above term back to (101) yields

$$V_2(x) = \max_{y_t} \int_x^{+\infty} \mathrm{E}\left[\int_\tau^{\tau+T} \left(y_t - \left(\frac{1}{2} + \frac{1}{\kappa}\right)\left(\frac{y_t}{\xi_t}\right)^2\right) dt\right] dF(\theta; x). \tag{104}$$

Point-wise maximization of (104) with respect to y_t yields the optimal output target y_t^* and effort e_t^*:

$$y_t^* = \gamma \xi_t^2 \tag{105}$$

$$e_t^* = \gamma \xi_t \tag{106}$$

where $\gamma = \kappa/(\kappa+2)$. Substituting (106) and (105) back into (98) yields the following information rent under the optimal contract:

$$R(\theta) = \int_x^\theta \mathrm{E}\left[\int_\tau^{\tau+T} \frac{1}{q}(\gamma\xi)^2 dt\right] dq = \int_x^\theta \gamma^2 q \mathrm{E}\left[\int_\tau^{\tau+T} \nu_t^2 dt\right] dq = \frac{\phi\gamma^2}{2}\left(\theta^2 - x^2\right) \tag{107}$$

Finally, substituting (105) back into (104) yields

$$V_2(x) = \int_x^{+\infty} \mathrm{E}\left[\int_\tau^{\tau+T} \frac{\gamma\xi_t^2}{2} dt\right] dF(\theta;x) = \int_x^{+\infty} \frac{\phi\gamma\kappa x^\kappa}{2}\theta^{1-\kappa} d\theta = \frac{\phi\gamma\kappa}{2(\kappa-2)}x^2. \tag{108}$$

Note that, similar to the baseline model, because the principal and the agent share the same discount rate (both 0), and there is no endogenous turnover during the production stage, all wage payments $\{w_t\}$ can be postponed until the end of the production period. Any W_{τ^-} carried over to the production stage can also be paid at the end of the production stage together with all the accrued wage payments.

C. The Search Stage

Let τ and $W_{t\leq\tau}$ denote the stopping time of the search stage (either due to progress to the next stage or contract termination) and the associated promised utility to the agent. By the martingale representation theorem for jump processes, given any investment strategy x of the principal, there exists a $\mathcal{F}_t$-predictable, integrable process β_t such that

$$dW_t = \beta_t(dN_t - \lambda(1 - F(x_t))dt). \tag{109}$$

Incentive compatibility of truthful reporting of θ requires that $W_\tau - W_{\tau^-} = R(\theta_\tau)$ if the contract moves to the next stage, which implies $\beta_t = \mathrm{E}[R(\theta_t)|\theta_t \geq x_t] = U(x_t)$ by the property of a martingale, where

$$U(x) \equiv \mathrm{E}\left[R(\theta)|\theta \geq x\right] = \int_x^{+\infty} \frac{\phi\gamma^2}{2}(\theta^2 - x^2)\left(\frac{\kappa x^\kappa}{\theta^{\kappa+1}}\right) d\theta = \frac{\phi\gamma^2 x^2}{\kappa-2}. \tag{110}$$

Therefore, under an incentive compatible contract with investment policy x_t,

$$dW_t = U(x_t)(dN_t - \lambda(1 - F(x_t))dt) \tag{111}$$

Then, Ito's lemma implies the principal's value function in the search stage solves the HJB equation:

$$0 = \max_x \quad -\delta - \lambda\left(\frac{\theta_{\min}}{x}\right)^\kappa U(x)V_1'(W) + \lambda\left(\frac{\theta_{\min}}{x}\right)^\kappa [V_2(x) - W - V_1(W)] \tag{112}$$

Substituting $U(x)$ from (110) and $V_2(x)$ from (36) into the HJB equation and rearrange terms yields:

$$V_1(W) = \max_x \quad \left(\frac{\phi\gamma\kappa}{2(\kappa-2)}\right)x^2 - \left(\frac{\phi\gamma^2}{\kappa-2}\right)x^2 V_1'(W) - W - \frac{\delta}{\lambda}x^\kappa\theta_{\min}^{-\kappa}. \tag{113}$$

The first order condition implies

$$\frac{\phi\gamma}{\kappa-2}\left[\kappa - 2\gamma V_1'(W)\right]x = \frac{\delta}{\lambda}\kappa x^{\kappa-1}\theta_{\min}^{-\kappa}. \tag{114}$$

The solution is

$$x(W) = \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right)V_1'(W)\right)\left(\frac{\lambda\phi\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}. \tag{115}$$

Substituting $x = \theta_{\min}$ into the above equation implies that there exists $\underline{W}$ such that

$$\theta_{\min} = x(\underline{W}) = \left[\left(1 - \left(\frac{2\gamma}{\kappa}\right)V_1'(\underline{W})\right)\left(\frac{\lambda\phi\gamma\theta_{\min}^{\kappa}}{\delta(\kappa-2)}\right)\right]^{\frac{1}{\kappa-2}}. \tag{116}$$

That is, $\underline{W}$ solves

$$V_1'(\underline{W}) = \left(1 - \frac{\delta(\kappa-2)}{\lambda\phi\gamma\theta_{\min}^2}\right)\frac{\kappa}{2\gamma}. \tag{117}$$

The rest of the proof is identical to that of Proposition 3 and is thus omitted.

References

Andrade, G., M. Mitchell, and E. Stafford (2001). New evidence and perspectives on mergers. *Journal of Economic Perspectives 15*(2), 103–120.

Antle, R. and J. Fellingham (1990). Resource rationing and organizational slack in a two-period model. *Journal of Accounting Research 28*(1), 1–24.

Arya, A., J. C. Fellingham, and R. A. Young (1994). Contract-based motivation for keeping records of a manager's reporting and budgeting history. *Management Science 40*(4), 484–495.

Azarmsa, E., L. Y. Liu, and S. Noh (2023). How does internal communication technology affect internal information: Theory and evidence. *Working paper*.

Baiman, S., M. Heinle, and R. Saouma (2013). Multistage capital budgeting with delayed consumption of slack. *Management Science 59*(4), 869–884.

Baldenius, T. (2003). Delegated investment decisions and private benefits of control. *The Accounting Review 78*(4), 909–930.

Baldenius, T., S. Dutta, and S. Reichelstein (2007). Cost allocations for capital budgeting decisions. *The Accounting Review 82*(4), 837–867.

Bastian-Johnson, N., T. Pfeiffer, and G. Schneider (2013). Multistage capital budgeting for shared investmentss. *Management Science 59*(5), 1213–1228.

Bastian-Johnson, N., T. Pfeiffer, and G. Schneider (2017). Two-stage capital budgeting, capital charge rates, and resource constraints. *Review of Accounting Studies 22*(2), 933–963.

Bebchuk, L., A. Cohen, and A. Ferrell (2009). What matters in corporate governance? *Review of Financial Studies 22*(2), 783–827.

Bergemann, D. and P. Strack (2015). Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory 159*, 819–853.

Bernardo, A., H. Cai, and J. Luo (2001). Capital budgeting and compensation with asymmetric information and moral hazard. *Journal of Financial Economics 61*(3), 311–344.

Bernstein, S., X. Giroud, and R. R. Townsend (2016). The impact of venture capital monitoring. *Journal of Finance 71*(4), 1591–1622.

Biais, B., T. Mariotti, G. 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, J.-C. Rochet, and S. Villeneuve (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica 78*(1), 73–118.

Billett, M. T., J. A. Garfinkel, and Y. Jiang (2011). The influence of governance on investment: Evidence from a hazard model. *Journal of Financial Economics 102*(3), 643–670.

Coles, J. L., N. D. Daniel, and L. Naveen (2014). Co-opted boards. *Review of Financial Studies 27*(6), 1751–1796.

Curello, G. and L. Sinander (2021). Screening for breakthroughs. Working paper, University of Bonn and Oxford University.

Dambra, M., O. Even-Tov, and K. George (2021). Should SPAC Forecasts be Sacked? *Available at SSRN 3933037*.

DeMarzo, P., M. Fishman, Z. He, and N. Wang (2012). Dynamic agency and the q theory of investment. *Journal of Finance 67*(6), 2295–2340.

DeMarzo, P. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance 61*(6), 2681–2724.

Drury, C. and M. Tayles (1996). Uk capital budgeting practices: some additional survey evidence. *The European Journal of Finance 2*(4), 371–388.

Dutta, S. and Q. Fan (2009). Hurdle rates and project development efforts. *The Accounting Review 84*(2), 405–432.

Dutta, S. and S. Reichelstein (2002). Controlling investment decisions: Depreciation-and capital charges. *Review of Accounting Studies 7*(2-3), 253–281.

Edmans, A., V. W. Fang, and E. Zur (2013). The effect of liquidity on governance. *Review of Financial Studies 26*(6), 1443–1482.

Eső, P. and B. Szentes (2007). Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies 74*(3), 705–731.

Fellingham, J. C. and R. A. Young (1990). The value of self-reported costs in repeated investment decisions. *The Accounting Review 65*(4), 837–856.

Feng, F. Z., T. Nohel, X. Tian, W. Wang, and Y. Wu (2023). The incentives of SPAC sponsors. *PBCSF-NIFR Working paper*.

Fich, E. M. and A. Shivdasani (2012). Are busy boards effective monitors? In *Corporate Governance*, pp. 221–258. Springer.

Franzoni, F. (2009). Underinvestment vs. overinvestment: Evidence from price reactions to pension contributions. *Journal of Financial Economics 92*(3), 491–518.

Gabaix, X. and A. Landier (2008). Why has CEO pay increased so much? *Quarterly Journal of Economics 123*(1), 49–100.

Gahng, M., J. R. Ritter, and D. Zhang (2022). SPACs. *Review of of Financial Studies*. Forthcoming.

Gershkov, A., B. Moldovanu, and P. Strack (2018). Revenue-maximizing mechanisms with strategic customers and unknown, markovian demand. *Management Science 64*(5), 2031–2046.

Glaeser, S. and M. Lang (2023). A review of the accounting literature on innovation. *Working Paper*.

Gompers, P., J. Ishii, and A. Metrick (2003). Corporate governance and equity prices. *Quarterly Journal of Economics 118*(1), 107–156.

Green, B. and C. R. Taylor (2016). Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects. *American Economic Review 106*(12), 3660–99.

Gregor, M. and B. Michaeli (2020). Board bias, information, and investment efficiency. *Working paper*.

Gregor, M. and B. Michaeli (2022). Board compensation and investment efficiency. *Working paper*.

Halac, M., N. Kartik, and Q. Liu (2016). Optimal contracts for experimentation. *Review of Economic Studies 83*(3), 1040–1091.

Harford, J. (1999). Corporate cash reserves and acquisitions. *Journal of Finance 54*(6), 1969–1997.

Heinle, M., N. Ross, and R. Saouma (2014). A Theory of Participative Budgeting. *The Accounting Review 89*(3), 1025–1050.

Hoffmann, F. and S. Pfeil (2010). Reward for luck in a dynamic agency model. *Review of Financial Studies 23*(9), 3329–3345.

Kirby, A. J., S. Reichelstein, P. K. Sen, and T.-Y. Paik (1991). Participation, slack, and budget-based performance evaluation. *Journal of Accounting Research 29*(1), 109–128.

Krähmer, D. and R. Strausz (2011). Optimal procurement contracts with pre-project planning. *Review of Economic Studies 78*(3), 1015–1041.

Lang, L. H., R. Stulz, and R. A. Walkling (1991). A test of the free cash flow hypothesis: The case of bidder returns. *Journal of financial economics 29*(2), 315–335.

Lin, C., F. Lu, R. Michaely, and S. Qin (2021). SPAC IPOs and sponsor network centrality. *Available at SSRN 3856181*.

Liu, B. and J. Lu (2018). Pairing provision price and default remedy: optimal two-stage procurement with private r&d efficiency. *RAND Journal of Economics 49*(3), 619–655.

Madsen, E. (2022). Designing deadlines. *American Economic Review 112*(3), 963–997.

Malenko, A. (2019). Optimal dynamic capital budgeting. *Review of Economic Studies 86*(4), 1747–1778.

Marinovic, I. and F. Varas (2019). CEO horizon, optimal pay duration, and the escalation of short-termism. *Journal of Finance 74*(4), 2011–2053.

Masulis, R. W., C. Wang, and F. Xie (2007). Corporate governance and acquirer returns. *the Journal of Finance 62*(4), 1851–1889.

Masulis, R. W., C. Wang, and F. Xie (2009). Agency problems at dual-class companies. *The Journal of Finance 64*(4), 1697–1727.

Mayer, S. (2022). Financing breakthroughs under failure risk. *Journal of Financial Economics 144*(3), 807–848.

McNichols, M. F. and S. R. Stubben (2008). Does earnings management affect firms' investment decisions? *The Accounting Review 83*(6), 1571–1603.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A Myersonian approach. *Econometrica 82*(2), 601–653.

Piskorski, T. and A. Tchistyi (2010). Optimal mortgage design. *Review of Financial Studies 23*(8), 3098–3140.

Poterba, J. M. and L. H. Summers (1995). A CEO survey of U.S. companies' time horizons and hurdle rates. *MIT Sloan Management Review* (Fall), 43–53.

Richardson, S. (2006). Over-investment of free cash flow. *Review of Accounting Studies 11*, 159–189.

Rozenbaum, O. (2019). Ebitda and managers' investment and leverage choices. *Contemporary Accounting Research 36*(1), 513–546.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies 75*(3), 957–984.

Varas, F. (2018). Managerial short-termism, turnover policy, and the dynamics of incentives. *Review of Financial Studies 31*(9), 3409–3451.

Williams, N. (2011). Persistent private information. *Econometrica 79*(4), 1233–1275.

Yin, C. and C. Zhu (2023). The value of activism: A hedge fund investor's perspective. *Working Paper*.