Private Equity Continuation Vehicles: A Model of Strategic Asset Transfers*

Victoria Ivashina
Harvard University
CEPR and NBER

Simon Mayer
Carnegie Mellon University

Ludovic Phalippou
University of Oxford

August 26, 2025

Abstract

We study private equity continuation vehicles (CVs), documenting their contractual features and developing a model that highlights the conflicts of interest they generate. While CVs can enhance efficiency by extending the holding period of high-potential assets, they also enable general partners (GPs) to extract rents by exploiting informational advantages and their intermediary position between legacy and new limited partners (LPs). In particular, GPs may transfer low-quality assets that should be liquidated, or high-quality assets that should be sold, thereby redistributing value across investor groups. The optimal legacy fund contract ties GP compensation to both deal- and fund-level outcomes and requires GP to invest more than the carry earned on the transaction. GP coinvestment mitigates adverse selection when low-quality assets are transferred, but strengthens incentives to move high-quality assets at below fair value.

Keywords: Private equity, Continuation vehicles, General partners, Limited partners, Agency conflicts, Asset transfer, Information asymmetry, Optimal contracts. JEL

Codes: G23, G24, G34, D82, D86

*We thank Sebastian Gryglewicz, Erwan Morellec, Gordon Phillips, and Mark Westerfield for their helpful comments. We are also very grateful to several market participants who generously shared data and insights on the continuation vehicle market.

In recent years, GP-led secondaries, or *continuation vehicles* (CVs), have become a central feature of the private equity (PE) landscape, reshaping traditional exit dynamics. In these transactions, general partners (GPs) sell a subset of existing portfolio companies to a newly structured fund they continue to manage. This contrasts with traditional exits, where companies are sold to other sponsors, strategic buyers, or through public offerings. CVs now account for a large share of both secondary market activity and overall PE exits. While they can create additional value and liquidity for limited partners (LPs), their rapid growth has also raised concerns about conflicts of interest, rent extraction, and governance.[1]

Despite their growing prominence, we still lack (i) systematic empirical evidence and facts on CVs and (ii) a conceptual framework for assessing their impact on fund economics, asset quality, and investor returns. In this paper, we address both gaps. First, we present key empirical facts about CVs. Second, we develop a framework that captures the multilayered conflicts of interest and moral hazard they entail. This framework shows how fund contracts should be structured to preserve the efficiency potential of CVs while limiting GP rent extraction. In addition, it yields predictions about when CVs arise, the types of assets they contain, and their distributional effects across investor groups.

In the first part of the paper, drawing on proprietary sources, we present institutional details and document key facts about the common structural elements of CVs. To the best of our knowledge, this evidence offers novel, granular insights into deal mechanics, investor participation, pricing, and contractual terms.

Continuation vehicles exhibit several distinctive contractual features. First, GPs are typically required to make sizable reinvestments, often exceeding 100% of the after-tax carry earned on the transaction. Second, carried interest is commonly structured through a tiered ratchet, with the full 20% allocation realized only once stringent hurdles such as a 20% gross IRR and a 2x MOIC are satisfied. Third, the investor base in CVs systematically differs from that of primary (legacy) funds, with secondary-market LPs playing a prominent role; this heterogeneity shapes bargaining outcomes and the pricing of transferred assets

In the second part of the paper, we develop a two-period model that aims to capture these stylized facts and to formalize the multilayered conflicts of interests between GPs and LPs. A GP initially manages a legacy fund investing in two portfolio companies. Success depends on the GP's costly, unobservable effort, creating a moral hazard problem addressed by performance-based pay (carried interest). At the end of the first period, the GP privately

[1] In what follows, we use common terms and their abbreviations interchangeably.

learns whether each firm in its portfolio has (1) succeeded (and can be exited), (2) failed, (3) or is "high potential," i.e., likely to succeed if given more time and effort. Crucially, only the continuation of high-potential projects is efficient. Failed ("zombie") projects or already successful ("crown-jewel") projects should be exited.[2]

Without CVs, the model ends at the end of the first period, leading to premature liquidation of high-potential assets. Allowing the GP to raise a CV changes this dynamic: the GP can transfer one asset to a continuation fund, extending ownership for another period. The CV raises capital from a mix of rolling legacy LPs and new LPs, who replace those exiting (due to liquidity needs). The GP coinvests both the carry earned from the legacy fund and personal capital, and a new contract governs incentives in the CV. Exiting LPs receive the transfer price of the continued asset, plus proceeds from the sale of the other asset.

In raising a continuation vehicle, the GP can exploit both its superior information about asset quality and its dual role as intermediary between legacy and new LPs. By selectively aligning with one group of investors through an opaque side arrangement, the GP is able to extract rents at the expense of others. These mechanisms generate two distinct types of distortions that lead to inefficient continuations and eventually harm investors.

First, zombie CVs arise when the GP overstates the quality of failed assets, offloading them to new LPs. Because exiting legacy LPs may benefit from such deals, the legacy fund contract may not deter them. Second, crown-jewel CVs arise when the GP understates the quality of successful assets to buy in cheaply, harming exiting LPs while benefiting new LPs who invest on favorable terms. The GP captures rents through its own coinvestment, side deals or kickbacks with new LPs, and its endogenous incentive compensation (rewards the GP for success without effort). Wary of being stripped of their best assets, legacy LPs may incentivize the GP to forgo such crown-jewel funds by increasing GP compensation. When this proves infeasible or inefficient, crown-jewel funds can arise, typically involving larger, higher-quality projects.

Overall, conflicts of interest arise between legacy LPs and continuation vehicle LPs. By exploiting its intermediary position between the two groups and selectively aligning with one at the expense of the other, the GP benefits from the low rollover rates among legacy LPs that we document. Higher rollover and coinvestment by legacy LPs ease these tensions and reduce agency rents and inefficiencies.

[2]In practice, the term "crown-jewel" is often used more broadly to include what we define as high-potential projects.

Notably, when all legacy LPs roll over their stake and no additional capital from new LPs is needed—a situation akin to extending the life of the legacy fund rather than raising a CV—the GP's agency rents are minimized, and the continuation decision is NPV-efficient (i.e., only high-potential projects are kept longer). Our analysis therefore identifies conflicts of interest between subsets of LPs, coupled with the GP's ability to exploit them. In practice, however, continuing solely with legacy LPs and not bringing in new investors is likely unattractive for the GP, whose rents are reduced, and for LPs with liquidity needs or a preference for timely exit.

We characterize the optimal continuation and legacy fund contracts that maximize LP value while limiting GP rent extraction. Normatively, this shows how contracts should be structured to address the multi-layered conflicts of interest between LPs and GPs. Positively, it highlights when and how different contractual designs can lead to the emergence of inefficient zombie and crown-jewel funds.

While continuation fund contracts are standard and pay the GP only if the continued asset succeeds, the optimal legacy fund contract is more nuanced, tying compensation to both deal- and fund-level performance, with terms contingent on whether a CV is raised. This contrasts with the outcome without CVs, where the GP is paid only if both projects succeed, i.e., compensation is fully based on fund performance. To discourage zombie CVs, GP's "crystallized" carry from the legacy fund should be reinvested in the CV, consistent with market practice.[3] To deter the GP from stripping the legacy fund of its best assets, legacy LPs reward the GP for individual deal success and increase GP compensation relative to the CV to make exiting successful firms within the legacy fund relatively attractive.

Besides characterizing the conflicts of interest inherent in CVs and the optimal fund contracts, we obtain four key results. First, the option to raise a CV always increases the GP's expected payoff, particularly when assets are riskier, more heterogeneous, and fewer legacy LPs roll over their stakes. Legacy LPs may gain or lose from CVs depending on whether efficiency gains outweigh rent extraction. Crucially, containing conflicts between LPs and the GP requires "ratcheting up" legacy-fund compensation relative both to CV compensation and to a no-CV benchmark. We show that if compensation is set too low, the GP has an incentive to raise crown-jewel funds at undervalued prices, to the detriment of exiting legacy LPs.

[3] "Crystallized carry" is a term commonly used in the PE industry to refer to the distributed carry. CV constitutes an exit for a legacy fund, thus the GP receives carry on the assets transferred to the CV.

Second, GP coinvestment deters zombie CVs and thus mitigates adverse selection at formation, but it can also encourage crown-jewel transfers at undervalued prices. Intuitively, the more the GP is expected to coinvest, the greater its benefit from acquiring a stake in the CV on favorable terms by transferring assets at undervalued prices. By contrast, legacy LP coinvestment and rollover reduce GP rent extraction and promote more efficient continuations, mitigating the emergence of both zombie and crown-jewel funds.

Third, CV asset quality is endogenous and may not reflect true value creation—crown-jewel funds, in particular, may include assets that should have been exited earlier. Such inefficient crown-jewel funds are more likely when fewer legacy LPs roll over, with GPs typically including larger assets in these funds.

Fourth, CVs are more likely during periods of high LP liquidity needs, which can increase the prevalence of both zombie and crown-jewel CVs. Such periods generally raise the volatility of asset quality while potentially improving its average level.

Our paper contributes to the nascent literature on secondary transactions. Nadauld, Sensoy, Vorkink, and Weisbach (2019) study the secondary market in which investors can buy and sell LP stakes in PE funds. Kastiel and Nili (2023) offer a legal perspective on continuation vehicles and discuss their institutional details. The insights that emerge from our paper are complementary to the work of Kastiel and Nili (2023).

We contribute to the theoretical literature on private equity and venture capital (e.g., Inderst, Mueller, and Münnich (2007); Axelson, Strömberg, and Weisbach (2009); Hochberg, Ljungqvist, and Vissing-Jørgensen (2014); Malenko and Malenko (2015); Gryglewicz and Mayer (2023); Maurin, Robinson, and Strömberg (2023)). Most closely related, Axelson et al. (2009) and Maurin et al. (2023) develop contracting theories of leveraged buyouts, focusing the GP's incentives to select projects over the fund's lifetime.[4] Our novel contribution lies in allowing for the endogenous exit or continuation of the fund's portfolio companies — capturing the economic essence of CV.

Our paper builds on Holmstrom and Tirole (1997), extending their framework to multiple assets, multiple periods, and adverse selection. Related are asset-sale models under asymmetric information (Leland and Pyle (1977) on a single asset and He (2009) on multiple assets), but our setting combines asymmetric information with moral hazard, introducing continuation funds as a novel sale mechanism. Unlike traditional asset sales, the asset is

[4]See Kaplan and Strömberg (2009) for an overview of leveraged buyouts, and Axelson, Jenkinson, Strömberg, and Weisbach (2013) for an empirical analysis of buyout leverage determinants.

partially sold to new LPs yet remains managed by the same GP under an optimal contract, creating conflicts of interest not previously studied.

The remainder of the paper proceeds as follows. Section 1 discusses the CV institutional details, and presents stylized facts. Section 2 presents the model. Section 3 derives the optimal contracts in the legacy and continuation funds. Section 4 examines implications for GP and LP behavior, asset quality, and market outcomes. Section 5 focuses on crown-jewel funds. Section 6 discusses extensions, including continuation funds of continuation funds. Section 7 concludes. All proofs are gathered in the Appendix.

1 Facts about Continuation Vehicles

The novel descriptive contribution of our paper is to provide structural details of a representative continuation vehicle transaction. While we also draw on multiple industry reports, these generally focus on volume trends in the GP-led secondary market and their composition, rather than structural features.

We have done several things to collect the information below. Our data come from two large secondary investors whose identities we keep confidential. One is a specialized secondary investor, and the other is a direct investments program of a large limited partner. First, we were able to review investment memos for 15 deals closed between Q1 2023 and Q2 2025 across the US and European markets, representing approximately $20 billion in equity.[5] This sample covers established mid- and large-cap private equity funds. Second, for both of our sources, we had access to internal summaries on GP compensation and expected performance of CVs for a broader set of deals spanning the period from Q1 2019 through Q1 2025. Finally, we also conducted interviews with investment teams to gather information about CV transactions they reviewed but did not participate in. In both instances, the investment teams believe that they have visibility into nearly the entire universe of CV transactions in the U.S. and Europe.

Currently, there is a wide variety of CVs on the market: a CV can combine multiple assets across different funds. More recently, this could even include assets from previous continuation vehicles. Based on our data, about half of the transactions are single-asset transactions. This is consistent with Evercore (2025) and Lazard (2025) reports, which

[5]Some of these transactions involve leverage, meaning that the actual deal volume is larger. For context, Evercore (2025) estimates that total GP-led deal volume in 2023 was $51 billion.

indicate that, in 2023 and 2024, single-asset transactions comprised close to 40% and 50% of the GP-led volume, respectively.

In what follows, we elaborate on the defining characteristics of CVs. As general background, our sources indicate that CVs remain more common among larger and more established GPs. Among smaller and younger PE firms, there is still a degree of stigma associated with pursuing a continuation vehicle. This underscores the role of a GP's reputation as an important mechanism for mitigating the conflicts inherent in continuation vehicles, although we do not model this point explicitly.

Fact 1 – CV purpose and timing. *Continuation vehicles represent a selective refinancing of portfolio assets that, on average, (i) extend the holding period for chosen assets by approximately four years beyond the original fund term, and (ii) raise additional capital to support and accelerate growth in those assets.*

In a typical GP-led secondary transaction, GP selects one or more assets from an existing, "legacy" fund(s) that it wishes to continue managing and transfers the asset(s) to a newly formed continuation vehicle created specifically to hold them. The transferred assets are typically considered "crown jewel" holdings—assets that GPs aim to retain and, importantly, continue growing beyond the standard fund timeline.[6] The market standard for CVs is a five-year life with two one-year extensions upon approval of the LP Advisory Committee (LPAC).[7] On average, the assets are transferred into the CV after they have been held for 5.1 years (median of 4.8 years) in the legacy fund, which is in line with the average holding period for PE portfolio companies. In the legacy fund structure, this typically occurs around year seven following the fund's closing.[8] Thus, CVs allow the GP to reset and extend the potential holding period for selected assets by approximately four years—seven years for the continuation fund, net of roughly three years remaining in the legacy fund. Figure 1 illustrates the typical timing of a continuation fund.

Given that the rationale for extending the holding period is often unrealized growth

[6] In line with the wide dispersion of returns across portfolio companies within the same fund—a defining characteristic of the private equity asset class (e.g., see Figure 2 in Braun, Jenkinson, and Schemmerl (2020) — there have historically been several instances of private equity firms retaining marquee assets across multiple fund vintages. Notable examples include FleetCor for Summit Partners and Visma for Hg Capital. The rise of continuation vehicles has, in many ways, formalized and streamlined this practice, offering a more accepted and structured approach to managing "crown jewel" assets over extended time horizons.

[7] A typical primary fund has eight-year life with two one-year extensions, i.e., ten years in total.

[8] A traditional private equity fund has maturity of ten years: eight years, plus two one-year extensions upon LP Advisory Committee approval.

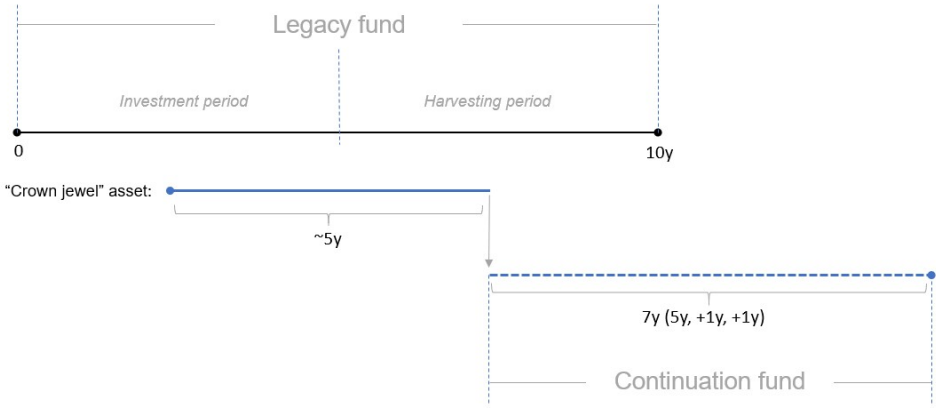


Figure 1: Continuation vehicle timing

potential, a typical CV also raises additional capital to support its investment thesis. In our sample, 15% of the CV equity exposure is unfunded on average (median of 11%). This additional capital is often intended to support growth through acquisitions.

In line with the narrative of remaining growth potential, CVs have an expected gross IRR of 20%. According to Lazard (2025), in 2024, single-asset CVs had an expected gross IRR of 21.3%. Multi-assets GP-led deals typically have a slightly lower target.

Fact 2 – Contractual LP-GP alignment in the continuation vehicle. *As a form of LP–GP alignment, continuation vehicles typically require the GP to coinvest 100% of the after-tax carry crystallized as a result of the transaction, along with additional out-of-pocket capital. The carry structure in the CV is usually governed by a multi-step ratchet, with the GP reaching the full 20% carry only after achieving at least a 20% gross IRR and a 2x MOIC. Management fees are calculated on invested capital and kept to a minimum.*

Effectively, creation of a continuation vehicle is a sale of an illiquid asset by legacy LPs to new LPs. The industry often refers to this practice as "selling to ourselves." Thus, an intrinsic feature of a continuation vehicle is the GP's dual role as both seller (on behalf of the

legacy fund) and buyer (on behalf of the CV), which gives rise to a conflict of interest. The central economic friction underlying this conflict is information asymmetry aggravated by the illiquid nature of the asset: GPs have superior insight into the asset's quality, trajectory, and risk profile. While assets in the primary funds are continuously valued by the third party, the marks tend to have a substantial lag. It is also well known that significant fraction of the value is realized at exit. Thus, asymmetry of information creates the potential for adverse selection, such as transferring underperforming ("zombie") or risky assets into CVs at inflated valuations. Alternatively, high-quality assets could be undervalued to secure favorable terms for the GP in the newly structured CV.

The economic incentives of a GP are complex. Since continuation vehicle is an exit for the legacy fund, the GP realizes carry when assets are passed to the CV.[9] At the same time, GPs economics are reset in a CV. (Table 1 provides a summary of typical terms in a continuation vehicle.)

We should also consider that selling LPs are investors that participate in the primary private equity market. According to Private Equity International, public pension funds accounted for 50.3% of the capital in private equity in 2024. Recognizing the complexity, potential conflicts of interest, and uneven expertise among participants in CV transactions, the U.S. Securities and Exchange Commission (SEC) adopted new rules in 2023 under the Investment Advisers Act aimed at GP-led secondaries. These rules mandate fairness opinions, enhanced disclosure of material conflicts, and increased transparency around valuation methodologies and carry crystallization mechanisms (SEC, 2023). The potential conflicts in the GP-led secondary transaction are also widely acknowledged in the industry (see, for instance, Persaud (2024)).

There are multiple implicit and explicit mechanisms that intend to mitigate the potential conflict of interest (i) between GP and legacy LPs, and (ii) between GP and the LPs in the continuation vehicle. We are particularly interested in contractual terms in the continuation vehicle which address GP and continuation vehicle LPs' alignment. In part, the point of the model presented in this paper is to evaluate effectiveness of these terms in ensuring alignment in the continuation vehicle and evaluate its implication for the legacy LPs, as well as to formalize these conflicts of interest.

Table 1 summarizes the key terms in a continuation vehicle and contrasts them with a

[9] As a rule, carry is in the money in CV transactions. In a European carry waterfall structure, carry is realized only after the preferred return on the entire called capital has been cleared. However, as noted earlier, CVs often occur around year seven of the fund—well into the fund's "harvesting stage".

primary fund. A few points are important to highlight:

- GPs coinvest in the CV by reinvesting 100% of the crystallized carry from the transferred assets and are typically required to contribute additional capital. There is a lot of variation in the GP contribution in excess of the carry roll (measured as % of the continuation fund assets under management). Similarly, the management teams of the portfolio companies are often expected to roll over most of their equity stake into the CV.

  Variation in a GP's total exposure to the CV is related to the size of the legacy fund: in smaller funds, GP exposure in the CV is typically around 10–11%, whereas in larger funds this percentage is substantially lower—even when the dollar contribution is actually greater in the latter case.

- There is variation in annual management fees but they tend to cluster around 50 basis points on the invested capital. Anchor investors typically receive an additional discount.

- It is common to use a carry ratchet in continuation vehicles. A typical carry structure has three tiers, with 20% carry only achieved upon realization of 20% gross IRR. It is also common to include MOIC hurdles in addition to IRR hurdles. For example, the highest carry bracket is typically conditional of 20% IRR and 2x MOIC.[10]

Fact 3 – CV process, buy-side. *The GP-led secondary market is highly competitive. In contrast to the primary fund market, LP investment demand for continuation vehicles is concentrated among a small number of large, specialized players.*

In a typical GP-led secondary, the GP engages a specialized advisory firm to run the process. Evercore, Lazard and Jefferies are among the most frequently used advisors in this market. The goal of the advisors is to ensure a fair and efficient process. Legacy LPs may sell their exposure at a negotiated price—typically based on third-party valuations or a competitive bidding process—or elect to "roll" their interests into the continuation vehicle. While the proportion of rolled capital varies across transactions, it is widely observed that most legacy investors choose not to roll over their stakes into the continuation vehicle. Our

[10] Although we have not encountered this in the transactions we have reviewed directly, it appears that some deals include so-called "super carry", which is an additional carry tier where the GP receives 25% carry once returns exceed a 25% gross IRR and a 3.0x MOIC.

sources indicated that only about 10% of legacy LPs by volume tend to roll their stake into a continuation vehicle. This number is also consistent with Kastiel and Nili (2023). Instead, demand for CVs is dominated by specialized secondary investors.

According to PEI Secondary Investor Ranking, in the five years between 2019 and 2023 the top-20 secondary firms raised close to $400 billion dollars, that is four times larger than the amount raised by the next 30 players. For comparison, according to Preqin, the total secondary deal volume over the same period was $480 billion.

The top continuation vehicle investors include specialized secondary funds affiliated with large alternative asset managers such as Blackstone, Ardian, HarbourVest, and Carlyle. The group also includes fund-of-funds platforms like StepStone, Pantheon, and Partners Group; dedicated secondary specialists such as Lexington Partners and Coller Capital; and traditional investment firms like Goldman Sachs. In addition, several large LPs pursue significant stakes in continuation vehicles through their direct private market programs.

Although the new buy-side CV investors often differ from those in the primary market, it is reasonable to assume that the legacy LPs—particularly those serving on the LPAC—are highly sophisticated, even if not specialized in secondary investments. In particular, many of these LPs have coinvestment programs, which suggests they have the capability to conduct independent due diligence on the portfolio companies.

Term	**Continuation Vehicle**	**Primary (Legacy) Fund**
Fund term	5 years, plus two 1-year extensions with LPAC approval	8 years, plus two 1-year extensions with LPAC approval
Funded	85%	0%
Management fee	0.25–0.75% on invested capital	1.5–2% of committed capital during investment stage, shift to invested capital base and step down in fees thereafter
Carried interest	10% over 10% gross IRR 15% over 15% gross IRR and 1.5–1.7x MOIC 20% over 20% gross IRR and 2x MOIC (100% GP catch up at all tiers)	20% above 8% gross IRR (100% GP catch up)
GP commitment	100% after tax carry roll, plus additional capital commitment	–

Table 1: Typical terms in a continuation vehicle vs. a primary fund

Fact 4 – CV process, information sharing. *As a market standard, all LPs in the legacy fund, as well as those in the continuation vehicle, receive the same information from the GP.*

There is an inherent information asymmetry between the GP and its LPs, this will be central to the model. To the extent that information asymmetry exists between exiting LPs and those in the continuation vehicle, it stems from institutional differences rather than from the flow of information.

A continuation vehicle transaction typically begins with securing an anchor investor—a buy-side investor (or group of investors) that promise to commit capital up to a specified amount. CV transactions are often bespoke and require significant due diligence and ongoing communication. As a result, outreach to legacy LPs usually does not begin until the transaction is sufficiently advanced and closure appears likely.

Once the GP has secured an anchor-investors for its continuation vehicle, all LPs in the legacy fund, as well as those in the continuation vehicle, receive the same information from the GP. Although GP-led secondaries are private transactions and may not be subject to the full scope of public securities regulations, GPs often follow market best practices to minimize legal risk and reduce investor friction. Additionally, under the Investment Advisers Act of 1940, GPs—acting as fiduciaries—are required to avoid conflicts of interest and ensure fair and transparent disclosure to LPs.

In U.S. public markets, information exchange is regulated by U.S. securities laws, including the Securities Exchange Act of 1934, Rule 13e-4, and Regulation 14E. These rules impose requirements around disclosure, equal treatment of holders, and minimum offer periods.

In practice, the GP typically establishes a secure "data room", an online repository where key documents related to the transaction are made available. These often include legal agreements related to the continuation vehicle, portfolio company performance data and financial projections, third-party valuation reports, and the GP's track record. All legacy LPs, advisors, and potential buyers are granted access to the same "data room."

Official information sharing is often complemented by informal guidance or direct communication ahead of the transaction launch. In general, an anticipated exit is communicated to LPs ahead of the event and tends to draw heightened attention, as this is when returns are realized. Exits involving a sale to a continuation vehicle typically attract even greater scrutiny due to the embedded conflicts of interest. As one GP noted, in the case of such an exit, they would speak with their largest LPs "three to five times before the LPAC meeting."

This is in line with Institutional Limited Partners Association (2019) continuation vehicles guidelines, which recommends that GPs should engage LPs "as early as possible in the process to provide the rationale for the restructuring and any alternatives considered."

Generally, it is reasonable to say that all LPs have access to the same "hard" information. However, as discussed earlier, the composition of legacy LPs differs substantially from the specialized investors who tend to dominate the CV side. Traditional LPs often face resource constraints and, in some cases, lack sophistication (see Ivashina and Lerner, 2019, for detailed examples). At the same time, legacy LPs benefit from a long history of both formal and informal communication with the GP regarding its portfolio investments. In sum, while it is difficult to broadly characterize whether legacy LPs or CV investors are better informed, it is plausible that in specific cases one group may hold an informational advantage (see, e.g., Ivashina and Lerner (2019)).

2 A Theory of Continuation Vehicles

We begin by outlining the mechanics of the model. In Subsection 2.2 we elaborate on the rationale behind the model setup.

2.1 Model Setup

We develop a model of private equity (PE) over three dates, $t = 0, 1, 2$. A risk-neutral general partner (GP) manages a PE fund on behalf of risk-neutral limited partners (LPs), with no discounting. At $t = 0$, the GP raises capital from LPs to manage a legacy (primary) fund, which operates until $t = 1$ and invests in two independent portfolio companies (henceforth, *projects*), indexed by $z \in \{A, B\}$.

Effort and Project Types. At $t = 0$, the GP exerts a hidden effort $e \in \{0, 1\}$, which determines the likelihood that project z succeeds by $t = 1$. With effort ($e = 1$), any project z is successful with probability $p_z \in (0, 1)$; without effort ($e = 0$), the success probability is zero. Not exerting effort yields a private benefit $\lambda^0 \geq 0$ to the GP. This moral hazard problem is standard in the spirit of Holmstrom and Tirole (1997); for simplicity, the same effort applies to both projects.

At the end of the first period (i.e., by time $t = 1$), each project is of one of three types, $s_z \in L, M, H$. A successful project (an H-type or *crown jewel*) can be sold to an external

buyer—such as another GP, a strategic acquirer, or via an IPO—at time $t = 1$ for a price of $X_z > 0$. Alternatively, it can be sold at time $t = 2$ for a potentially higher price of $X_z + x$, where $x \geq 0$ captures the additional value from continuing an already-successful project.

If project z is not successful by time $t = 1$, it cannot be sold at price X_z and its time-1 value to external buyers is normalized to zero. It then falls into one of two types:

- *High-potential* (M-type), with probability $1 - \delta_z$, where $\delta_z \in [0, 1]$ is exogenous;[11]
- *Zombie* or *failed* (L-type), with probability δ_z, i.e., δ_z is the failure probability.

A *zombie* project has permanently zero value to external buyers.

On the other hand, a *high-potential* project may succeed if continued to $t = 2$ and the GP exerts effort again, i.e., $e_z = 1$ at $t = 1$. In that case (as in the first period), project z succeeds with probability p_z, yielding a payoff of X_z at $t = 2$ (the price at which it can be sold to external buyers); otherwise, it yields zero. Low effort, $e_z = 0$, provides the GP with private benefits $\lambda_z \geq 0$ (where λ_z may differ from λ^0) but reduces the probability of success to zero.

Continuing any type of project requires investment of $I > 0$ dollars. We assume that only the continuation of an M-type project (under GP effort) is efficient, yielding positive NPV of $p_z X_z - I$, while continuing an H- or L-type project is inefficient, yielding negative NPV relative to exit at time $t = 1$. Formally:

$$p_z X_z > I > \lambda_z + x.$$

Continuation Vehicle. At time $t = 1$, the GP privately observes the type or state $s_z \in \{L, M, H\}$ for each project $z \in \{A, B\}$. The GP then decides whether to (i) sell the project to an external buyer at the given price, or (ii) transfer it to a continuation vehicle, which extends ownership to time $t = 2$ and is managed by the same GP.

We impose the following simplifying assumptions, consistent with observed practice. First, at most one project can be transferred to the CV, reflecting that CVs typically contain a single firm and are smaller than legacy funds. Second, a CV is initiated only if at least one project has succeeded by time $t = 1$. Intuitively, the GP seeks to avoid the reputational damage of poor performance across both vehicles and therefore raises a CV only if it knows that at least one project is successful.

[11] For M-type projects, a $t = 1$ price of zero does not mean the project is worthless, but rather that it requires additional GP involvement to realize its potential, or that external buyers do not value its potential.

If firm z is transferred to the CV, the investment and endogenous time-1 payout to legacy LPs driven by liquidity needs are financed by three sources: (i) the GP, (ii) legacy LPs participating in the continuation vehicle, and (iii) new LPs arriving at $t = 1$. New LPs are competitive, risk-neutral, and have deep pockets. The GP coinvests an exogenous amount $\kappa \geq 0$, which it is endowed with at the time of CV formation. In exchange, the GP receives a stake in the continuation vehicle.[12]

A fraction $\xi \in [0, 1]$ of legacy LPs (normalized to a unit mass) invest ξI in the continuation fund (i.e., I dollars per rolling LP) and "roll" into the CV. (The low observed rollover rates correspond to a small value of ξ.) The remaining fraction $(1-\xi)$ exits and receives a payout of $(1-\xi)\ell_z$ dollars to meet liquidity needs, where the payout per exiting LP $\ell_z \geq 0$ is determined endogenously (as characterized later). In aggregate, legacy LPs coinvest $\xi I - (1-\xi)\ell_z$ dollars. Thus, legacy LPs' coinvestment is limited by liquidity needs and capital constraints.

Consequently, new LPs invest the following amount of capital in the CV:

$$\text{Inv}_{\text{New}} := \underbrace{I}_{\substack{\text{Investment} \\ \text{cost}}} - \underbrace{\kappa}_{\substack{\text{GP} \\ \text{coinvestment}}} - \underbrace{\xi I}_{\substack{\text{Legacy LP} \\ \text{coinvestment}}} + \underbrace{(1-\xi)\ell_z}_{\substack{\text{Time-1 payout} \\ \text{(Legacy LP liquidity needs)}}} . \tag{1}$$

New LPs are competitive (consistent with Fact 3) and therefore break even, given their beliefs, as will be made precise later.

Contracts, Incentives, and Information Sets. At time $t = 0$, legacy LPs sign a contract $\mathcal{C}^{Leg}$ with the GP, which specifies compensation based on $t = 1$ cash flows (from selling firms to external buyers or the CV).

If project z is transferred to a CV, the continuation fund LPs sign a separate contract $\mathcal{C}_z^{CV}$ with the GP, specifying compensation contingent on cash flows at $t = 2$.

We focus on incentive-compatible contracts $\mathcal{C}^{Leg}$ and $\mathcal{C}_z^{CV}$ that (i) induce effort in both periods and (ii) in the case of the legacy fund contract, induce truthful reporting of project types s_z to legacy LPs. Payments to the GP must be non-negative (reflecting limited liability), and contracts are determined to maximize the LP value in the respective fund.

As is common in private equity, the GP can always engage in private side contracts with subsets of LPs, meaning that overall contractual arrangements are private and observable or verifiable only to the parties involved. Consequently, new LPs cannot observe or verify

[12] We preview that, in addition to the exogenous coinvestment of capital κ, there will also be an endogenous coinvestment, funded from a portion of the GP's compensation received from the legacy fund.

the full contractual terms or communications between the GP and legacy LPs. Likewise, at time $t = 1$, when the CV is formed, the GP and new LPs may privately communicate and strike side deals that remain unobserved by legacy LPs.

Timing. The timing of events is as follows:

1) At $t = 0$, LPs invest in the legacy fund and sign contract $\mathcal{C}^{Leg}$ with the GP. The GP chooses effort $e \in \{0, 1\}$.

2) At $t = 1$, the GP privately observes project types $s_z \in \{L, M, H\}$. The GP reports project types to legacy LPs and decides whether to sell or continue a project.

3) If a CV (with firm z) is formed at $t = 1$, new LPs invest at endogenous terms. The contract $\mathcal{C}_z^{CV}$ is signed between GP and the investors in the continuation vehicle, and the GP exerts effort $e_z \in \{0, 1\}$.

4) Cash flows from exited projects are realized and legacy fund closes. The GP is compensated per $\mathcal{C}^{Leg}$. Date $t = 1$ ends.

5) At $t = 2$, if a project was continued, cash flows from the continuation fund are realized. The GP is compensated per $\mathcal{C}_z^{CV}$.

The model unfolds as a sequence of three subgames: (i) the legacy fund contracting stage at $t = 0$, (ii) the continuation decision and fund formation at $t = 1$, and (iii) the CV contracting stage, also at $t = 1$. Rather than presenting the entire game upfront, we solve these subgames sequentially, proceeding backward to characterize the subgame perfect equilibrium.

2.2 Discussion of Assumptions

This section provides additional discussion and motivation for key model assumptions. It does not introduce any new formal elements, but offers more background and justification of the assumptions.

Overall Structure. In our model, the legacy fund holds two projects—the minimal number needed to constitute a portfolio. The GP can choose between these two projects for the continuation fund, adding an extra layer of optionality: in addition to deciding whether to raise a CV, the GP can also select which firm to include in the CV. This feature requires

modeling a portfolio and, importantly, distinguishes our setting from one in which the legacy PE fund can only continue a single asset.

We assume that at CV formation, a project is one of three types: L-type, M-type, or H-type. Three types are the minimum needed to capture that continuation can be inefficient either because a project has failed (underperformed) or because it is already successful and requires no further input. Both situations are relevant in the context of continuation funds, corresponding to the "zombie" and "crown-jewel" cases.

In our framework, crown-jewel and zombie CVs are inefficient by assumption and destroy value for certain investor groups, whereas high-potential CVs create value. We note, however, that in practice the term "crown-jewel" (and "crown-jewel CV") is sometimes used more broadly to include what we define as high-potential projects.

PE Value Creation through GP Effort. We assume that value creation in PE depends on the GP's costly, unobservable effort, creating a standard moral hazard problem (e.g., Casamatta, 2003; Schmidt, 2003; Gryglewicz and Mayer, 2023; Maurin et al., 2023). In our setting, this moral hazard generates endogenous incentive-based pay, such as carried interest, which shapes the GP's trade-off between exiting investments and transferring them to a CV. Effort reflects the GP's role in monitoring, board activism, strategic guidance, and designing managerial incentives, and carries an opportunity cost.

Contracts. GPs and investors can engage in private side deals or contractual arrangements that are not observable by other LPs. This is true for the legacy fund and for the continuation vehicle. This reflects the idea that contract details are often private to the signatory parties—even if a PE fund is governed by a common limited partnership agreement (LPA), the GP and the participating LPs can enter into side arrangements ("side letters") that render the overall contractual relationship effectively private to outsiders and alter the overall incentive structure.

In our model, this has two implications. First, new LPs may not fully observe the legacy contract. This opacity weakens the ability of the GP and legacy LPs to credibly commit not to transfer zombie firms to new LPs. Second, exiting LPs cannot observe the final contractual terms between the GP and new LPs.. In our model, the GP receives an explicit or implicit reward from new LPs when the CV outperforms, because the GP has transferred a successful asset at an undervalued price, boosting the new LPs' returns. Such rewards may also take the form of kickbacks to the GP—for example, when new LPs invest in a follow-up fund managed by the same GP, thereby rewarding the GP with future commitments.

In this context, we assume that an H-type project, if continued to $t = 2$, yields a higher cash flow $X_z + x$ with $x \geq 0$, even though continuation is inefficient in NPV terms. This captures the notion of overperformance: when an H-type is continued, the CV outperforms relative to an M-type. Such overperformance is observable and enables the GP to strike a side deal with new LPs that rewards it. Crucially, one could relax the modeling of strict overperformance, and also allow M-type projects to realize cash flow $X_z + x$ with some lower probability, without altering our findings.[13]

Here, x represents the incremental value increase of an already successful asset. This value is not generated through GP effort and thus does not constitute value creation—unlike the value creation that moves a project into state $s_z = H$. By assumption, holding an asset longer to realize this incremental value x is inefficient.

CV Assumptions. We assume that a fraction of legacy LPs must exit for liquidity reasons, creating a natural role for new LPs to replace them in the CV. While simplistic, this assumption reflects the rigidity of LP mandates and capital constraints in practice. Even LPs without immediate liquidity needs may be unable to roll over due to legal, regulatory, or institutional restrictions.

For simplicity, we assume that liquidity needs are i.i.d. shocks to each member of a unit mass of legacy LPs, so that the overall fraction of LPs experiencing a liquidity need is deterministic. Thus, unlike in Maurin et al. (2023), we do not model aggregate liquidity needs, though we later discuss such situations through comparative statics in ξ.

For tractability, we also assume that continuation occurs only if at least one project is successful by the interim date, reflecting reputational concerns—GPs rarely raise CVs when their entire portfolio underperforms. We also allow only one firm to be continued, consistent with market practice where single-asset continuation vehicles dominate. While neither assumption is central to the mechanisms we study and both could be relaxed, they simplify the analysis by reducing the number of possible cases at CV formation.

Continuation requires renewed GP effort and a reinvestment cost, capturing the need for fresh capital often seen in follow-ons, add-ons, or growth initiatives. More broadly, this reflects that continuation typically involves at least some additional resources, even if modest. Finally, we omit an explicit investment cost for the initial legacy fund; this is innocuous, as such a cost would simply reduce initial payoffs by a constant without affecting contract

[13] In the Appendix B.4, we show that this would not affect the GP's ability to extract rents through a side deal with new LPs.

design or continuation decisions.

Optimal Contracting. We highlight the multi-layered conflicts of interest between LPs and GPs and characterize the optimal contracts. Our approach has both normative and positive implications. Normatively, it shows how financing and contracts should be structured to maximize LP value and limit GP rent extraction in CVs, leading to features such as compensation tied to both fund- and deal-level performance. These features are often observed in practice, though they may also arise even without explicit anticipation of CVs. Positively, it explains how and when different contractual arrangements can give rise to inefficient zombie and crown-jewel funds. In either case, our analysis clarifies how contractual provisions address these conflicts.

In practice, although primary market LPAs have remained largely unchanged with the rise of CVs, informal norms have developed—for example, that transfers should occur only after achieving at least a 2x gross MOIC, "ideally 3x." In our framework, this practice is reflected in two ways: (i) CVs arise only when the legacy fund has performed sufficiently well (consistent with our assumption that a CV can be raised only if at least one project succeeds), and (ii) GP payoffs—whether through direct compensation or implicit rewards from LPs (e.g., future coinvestments)—are partly contingent on CV exits. More broadly, this suggests that legacy LPs anticipate CV exits when committing capital and engaging with GPs, consistent with our model.

3 Solution

3.1 Continuation Fund: Contracts and Financing Arrangements

When project z is continued, the LPs and the GP sign a contract $\mathcal{C}^{CV}$, which stipulates payments to the GP contingent on $t = 2$ cash flows. When $s_z = M$, meaning project z is high-potential, it either generates a cash flow of zero—in which case the GP naturally receives no payment—or a cash flow of X_z, in which case the GP is paid $w_2(z)$. Thus, the GP exerts effort if and only if:

$$p_z w_2(z) \geq \lambda_z \tag{IC-C}$$

This is a standard incentive constraint to ensure that the GP exerts a costly effort, which requires incentive-based compensation.

The GP contributes κ dollars of its own capital to the CV. We assume that, when $s_z = M$,

the GP breaks even within the continuation fund:[14]

$$p_z w_2(z) \geq \kappa. \tag{PC-C}$$

Because continuing an L- or H-type project is inefficient, including such a project in a CV inevitably harms at least one group of LPs who must bear the resulting costs. Accordingly, whenever the GP opts for continuation, it publicly presents the project as type M. To support this claim, whether true or not, the GP offers a public contract $\mathcal{C}_z^{CV}$ designed to maximize the value of the LPs as if $s_z = M$, i.e., to minimize agency rents subject to incentive and participation constraints. The participation constraint (PC-C) requires the GP to coinvest at fair terms, given its claim that $s_z = M$ and its effort choice $e_z = 1$. In practice, this reflects that GPs typically acquire CV shares at a fair valuation—where, in our context, the valuation is computed under the GP's claim that $s_z = M$ for consistency.[15]

We impose a standard monotonicity condition: if the project ultimately generates $X_z + x$, the GP must still receive at least $w_2(z)$—the payment tied to the lower X_z outcome.[16] We already impose that this constraint binds, i.e., a deviating GP prefers to preserve the M-type claim rather than reveal the deviation. This yields:

Proposition 1. *If project z is continued, the continuation fund contract $\mathcal{C}_z^{CV}$ requires the GP to coinvest κ and pays the GP $w_2(z)$ whenever the project can be sold at time $t = 2$ for at least X_z. The payment satisfies $p_z w_2(z) = \max\{\kappa, \lambda_z\}$.*

Note that our modeling of CV contract determination is not overly restrictive: although the CV is governed by the public contract $\mathcal{C}_z^{CV}$, signed under the GPs claim that $s_z = M$, the GP may still arrange private side deals with either legacy or new LPs, thereby altering the effective incentive structure.

[14] The GP is exogenously endowed with κ dollars at $t = 1$, which is not available at $t = 0$. This condition effectively treats the GP's outside option as zero, abstracting from any payoff under the legacy fund contract, which is neither observable nor verifiable to new LPs.

[15] As discussed in the next section, new LPs setting the price for CV shares may hold beliefs different from $s_z = M$ with probability one. See Appendix B.3 for an alternative formulation of the break-even constraint based on new LPs' beliefs about the prevailing state.

[16] Note that a $t = 2$ cash flow of $X_z + x$ reveals the continued project as H-type and thus a potential misrepresentation. We could easily allow that also an M-type project may produce cash flow $X_z + x$: as long as the H-type project is more likely to do so, our findings remain unchanged.

3.2 Continuation Fund Formation

Suppose the GP considers continuing project z, while the other project — denoted by $-z$ — has succeeded. Recall that we focus on incentive-compatible legacy fund contracts $\mathcal{C}^{Leg}$, which require the GP to truthfully reveal the state s_z to the legacy LPs. In this subsection, we therefore take truthful reporting of s_z as given, while the full characterization of the legacy fund contract under incentive-compatibility conditions is provided later.

Given s_z, the GP and legacy LPs jointly decide whether to transfer a project and how much capital to raise from new LPs, effectively acting as a coalition. Thus, the contract $\mathcal{C}^{Leg}$ also induces a continuation strategy that maximizes joint surplus, subject to incentive compatibility constraints. When the GP reports $s_z = H$ to legacy LPs, both parties jointly maximize their surplus by exiting the project, as there are no gains from trade and continuation is inefficient. Consequently, in equilibrium, only L- and M-type projects are continued — this changes in Section 5 where H-type projects may be continued as well in equilibrium (after we slightly modify our baseline assumptions).

Financing. The continuation fund is fully equity-financed, with a unit mass of shares. Each share pays $X_z - w_2(z)$ in the event of success and zero otherwise. That is, we define equity net of GP compensation; the GP also has an equity-like stake in the CV via $w_2(z)$.[17] The endogenous equity value — gross of the investment I — equals V_z, so the dollar return on equity equals $V_z - I$.

Let θ_z denote the fraction of continuation fund shares retained by the existing LPs, so new LPs buy $1-\theta_z$. According to (1), the GP raises $(1-\xi)(\ell_z + I) - \kappa$ dollars from new LPs to cover both the investment cost and the payouts of legacy LPs driven by liquidity needs. We model these liquidity needs by requiring that legacy LPs' payout at time $t = 1$ exceeds a minimum threshold.

Assumption 1 (Liquidity Needs)**.** *At $t = 1$, payouts to legacy LPs must satisfy $\ell_z \geq V_z - I$.*

The interpretation is as follows: there is a unit mass of legacy LPs, of which a fraction ξ rolls into the CV, while the remaining fraction $1 - \xi$ exits by selling their shares at the fair value $V_z - I$ (net of investment cost). Each rolling LP contributes capital I in exchange for continuation fund shares worth V_z gross of investment.[18]

[17] In equilibrium, only L- and M-type projects are continued. These projects generate binary cash flows and the GP is compensated only for success. Thus, the GP is effectively an equity holder in the continuation fund.

[18] As will become clear later, each rolling LP receives at least one share per I dollars coinvested, resulting in

A. Zombie Funds: Legacy LPs and the GP Benefit Off New LPs The contractual arrangements and communication between legacy LPs and the GP are neither observable nor verifiable to new LPs — for instance, because the parties can privately renegotiate contract terms or enter into a side deal. New LPs cannot observe or verify the state s_z and the GP's report to legacy LPs. Thus, the coalition of legacy LPs and the GP cannot credibly commit to excluding failed projects from continuation—they may benefit at the expense of new LPs by raising a "zombie fund."[19]

The continuation decision and equity sale form a signaling subgame with *adverse selection*: the GP–legacy LP coalition knows $s_z \in \{L, M\}$ (recall that H-type projects are not continued in equilibrium), while new LPs do not. Given s_z, the coalition decides (i) whether to continue and transfer project z and (ii) how much capital $(1-\xi)(\ell_z + I) - \kappa$ to raise from new LPs by choosing $\ell_z \geq V_z - I$. Observing these, new LPs form their beliefs about the transferred firm's type and invest at price V_z which is fair given their beliefs. Specifically, new LPs are competitive and merely break even under price V_z.

The following Proposition — proven in Appendix A.1 — summarizes the equilibrium.

Proposition 2 (Adverse Selection at CV Formation)**.** *The continuation-stage subgame (for firm z) admits a unique Perfect Bayesian Nash equilibrium, using the intuitive criterion as the refinement. In equilibrium, the GP always continues an M-type firm and never continues an H-type firm. The GP continues an L-type firm—mimicking type M—with probability γ_z (as characterized in (A.8).) Moreover, the following holds:*

1. *New LPs buy fraction $1 - \theta_z$ of continuation fund equity at price V_z, while legacy LPs retain fraction θ_z, where*

$$\theta_z = \min\left\{\xi + \frac{\kappa}{V_z}, 1\right\} \quad \text{and} \quad V_z = \frac{(1-\delta_z)p_z}{(1-\delta_z) + \delta_z\gamma_z}\,(X_z - w_2(z)). \tag{2}$$

2. *The GP coinvests κ dollars and legacy LPs and the GP raise the minimum amount of external financing to continue, i.e., $\max\{0, (1-\xi)(\ell_z+I)-\kappa\}$ dollars where $\ell_z = V_z - I$.*

3. *Coinvestment by legacy LPs or the GP in the continuation fund mitigates adverse selection: as κ or ξ increases, θ_z increases and γ_z decreases.*

a dollar return from coinvesting that is (weakly) greater than $V_z - I$. One could also change the assumption on liquidity needs (e.g., $\ell_z \geq V_z$) without affecting the model's key insights.

[19] One might worry that selling a failed project harms rolling legacy LPs while benefiting those who exit. However, if continuation increases the total surplus for the GP and legacy LPs, (hidden) side transfers can make all coalition members — including rolling LPs — better off than under liquidation.

Thus, the equilibrium can be (partial) pooling ($\gamma_z > 0$) or separating (γ_z). When coinvestments by the GP (κ) or legacy LPs (ξ) are larger, γ_z declines, marking a shift toward separating equilibrium.

B. Crown-Jewel Funds: New LPs and the GP Benefit Off Legacy LPs. The GP and new LPs can profit at the expense of exiting legacy LPs by raising a so-called "crown-jewel fund"—that is, by transferring an already successful project at an undervalued price. In equilibrium, only L- and M-type projects are continued at price V_z. However, the GP may deviate by continuing an H-type project while misrepresenting it as type M and setting the fair transfer price under this false claim, V_z. Then, new LPs acquire a fraction $1-\theta_z$ of the continuation fund equity at price V_z, paying $(1-\theta_z)V_z$ to exiting legacy LPs. In reality, each share is worth $X_z + x - w_2(z)$, which exceeds V_z. As a result, continuing an H-type project, while understating its quality as M-type, generates a total "surplus" of

$$\Omega(z) := (1-\theta_z)\big(X_z + x - w_2(z) - V_z\big) \tag{3}$$

for the GP and the new LPs at the expense of legacy LPs, who exit at undervalued price.

Crucially, the GP may extract part of this surplus by entering into a private side deal with new LPs. Specifically, upon observing that the project is H-type at time $t = 1$, the GP can offer a side contract that specifies a payment of $\widetilde{w}_2(z)$ to the GP if the project ultimately generates cash flows of $X_z + x > X_z$, and zero otherwise.[20] This side contract reflects private arrangements in which the GP is rewarded by new LPs—explicitly or implicitly—when the CV outperforms the price they paid, generating superior returns. In practice, such rewards may take the form of future investments by new LPs in the GP's funds, informal kickbacks, or formal agreements.

We denote by $\alpha \in \{0, 1\}$ whether the GP has full bargaining power in offering the side deal. When $\alpha = 1$, the GP can propose a contract under which new LPs merely break even, allowing the GP to extract the entire surplus $\Omega(z)$ from continuing an H-type project. When $\alpha = 0$, the GP either lacks bargaining power or the side deal is infeasible (e.g., when $x = 0$ or for exogenous reasons), so the GP captures no surplus and it remains with the new LPs in the event of deviation.

[20] This contract would not pay anything if the GP had an L- or M-type project. Thus, for new LPs, there is no cost in entering the contract, when the project is not H-type. If we assumed an M-type project could also produce cash flows $X_z + x$, albeit at lower probability, the GP could still extract rents from new LPs, as we show in Appendix B.4.

Thus, by the private side deal, the GP can extract in total $\alpha\Omega(z)$. Note that $\Omega(z)$ is fully characterized by the quantities from Propositions 1 and 2. The extent to which the GP can extract rents by transferring successful projects at undervalued price is larger when fewer legacy LPs roll, i.e., θ_z is lower, or the transfer price V_z is lower.

3.3 Legacy Fund Contract

At $t = 0$, legacy LPs and the GP sign a contract $\mathcal{C}^{Leg}$, which stipulates payments contingent on (i) the observable cash flows at $t = 1$, (ii) the reported state, and (iii) the continuation decision and terms. These payments are made when cash flows at $t = 1$ are realized—that is, after the continuation decision. We focus on incentive-compatible contracts that induce (1) full effort, (2) truthful reporting of s_z toward legacy LPs, and (3) a continuation decision by the GP consistent with the equilibrium from the continuation stage subgame in Proposition 2. Project z is continued only if the other project $-z$ was successful; successful projects are always exited at $t = 1$. While M-types are always continued, L-types are continued with probability γ_z.[21] The following Proposition characterizes the optimal, incentive-compatible contract that maximizes legacy LPs' payoff. Quantities $(V_z, \Omega_z, \gamma_z)$ are determined in the continuation-stage subgame from the previous section and are thus taken as given. To avoid edge cases and unnecessary case distinctions, we restrict attention to parameter values for which the payment expressions below are positive.[22]

At $t = 1$, four cases arise. If project z is successfully exited while the other is continued, the GP receives $w_{HM}(z)$ (project z in state H, with $-z$ reported as M). If project z is successfully exited while the other is exited at a low price, the GP receives $w_{HL}(z)$ (project $-z$ reported as L). If both projects are successfully exited, $(s_A, s_B) = (H, H)$, the GP receives w_{HH}. If both are exited at a low price, the GP receives nothing. Here, the subscripts indicate the reported states, while the argument (z) denotes the project that succeeds when only one does. The optimal payments in the first three cases are determined as follows:

[21] Legacy LPs and the GP cannot commit to excluding zombies. The only possible gain from committing to $\gamma_z = 0$ is a higher continuation price V_z. But since new LPs cannot observe verify legacy contract terms, the GP and legacy LPs can renegotiate to offload a failed project if it increases surplus, captured by the continuation stage game.

[22] If an expression were negative, the corresponding payment would simply be set to zero, without affecting the main message. A sufficient condition is $\frac{\lambda_z(1+p_z)}{p_z} \geq \lambda_{-z}$ for all $z, -z$, ensuring $w_{HM}(z)$ in (A.17) is positive. This mild condition holds, for instance, for symmetric projects.

Proposition 3 (Legacy Fund Contract)**.** *The optimal legacy contract compensates the GP if and only if at least one project is successfully exited at time $t = 1$, with:*

1. *If the GP successfully exits project z while continuing project $-z$, the GP is paid*

$$w_{HM}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}. \quad (4)$$

2. *If the GP successfully exits project z while exiting $-z$ at a price zero, the GP is paid*

$$w_{HL}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa. \quad (5)$$

3. *If the GP exits both project successfully, the GP is paid*

$$w_{HH} = \max_{z=A,B} \left\{ w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z), w_{HH}^{IC} \right\}, \quad (6)$$

where w_{HH}^{IC} is defined in the Appendix in (A.14) and satisfies $w_{HH}^{IC} \leq \frac{\lambda^0}{p_A p_B}$.

In the main text, we present only the essential equilibrium quantities; closed-form expressions for payoffs are given in the Appendix: NPV in (A.23), GP payoff Π_{GP} in (A.21), and LP payoff Π_{LP} in (A.24).

The intuition behind this proposition is as following: First, the GP is paid $w_{HM}(z) > 0$ in the legacy fund after successfully exiting project z while continuing project $-z$. The payment $w_{HM}(z)$ is necessary to deter strategic behavior, whereby the GP shifts crown-jewel (H-type) projects into the continuation fund at undervalued prices while leaving underperforming ones behind in the legacy fund.

To see how (4) is determined, consider the case where project z is successful and project $-z$ remains a high-potential M-type. Legacy LPs incentivize the GP to continue $-z$ while exiting z. By doing so, the GP receives $w_{HM}(z)$ from successfully exiting z, along with an expected payoff of $p_{-z}w_2(-z)$ from the continuation fund, requiring a coinvestment of κ. However, the GP may deviate by continuing the crown-jewel z — while misrepresenting it as an M-type — and exiting $-z$ at price zero. The GP can extract $\alpha\Omega(z)$ through the side deal with new LPs, receive $w_2(z)$ with certainty at $t = 2$, and enjoy private benefits λ_z from shirking, while still investing κ. If legacy LPs wish the GP to continue $-z$ even after failure (i.e., $\gamma_{-z} > 0$), incentivizing continuation of $-z$ becomes more difficult, and may require an additional payment of $(p_{-z}w_2(-z) - \lambda_{-z})$, which is zero for $\lambda_{-z} \geq \kappa$.

Second, the GP must also be compensated—receiving $w_{HL}(z) > 0$—when both projects

are exited at $t = 1$, but only project z has succeeded. This deters the GP from rolling over the successful project into the continuation fund. More broadly, the contract must reward even a single success; otherwise, the GP may shift the crown-jewel to the continuation fund while leaving the other project behind.

Third, suppose both projects succeed. Legacy LPs incentivize the GP to exit both projects through the payment w_{HH}. However, the GP may deviate by continuing one project z and pretending it is of type $s_z = M$, while exiting $-z$. In that case, she receives $w_{HM}(-z)$ from exiting $-z$ at $t = 1$, extracts surplus $\alpha\Omega(z)$, earns $w_2(z)$ at $t = 2$, and enjoys private benefits λ_z, while investing κ. Crucially, the GP can choose which of the two successful projects to continue, allowing her to extract greater rents. This reflects the optionality in raising a continuation fund—not only deciding whether to raise one, but also which assets to transfer—thus increasing the compensation needed for incentive compatibility.

Fourth, there is a standard time-0 incentive compatibility constraint of the form $w_{HH} \geq w_{HH}^{IC}$, where w_{HH}^{IC} depends on λ^0, the private benefits of shirking. This constraint binds only when λ^0 is sufficiently large. Since it is neither novel nor specific to continuation funds, we do not focus on it. Instead, we assume the following:

Assumption 2. *Assume that λ^0 is not too large, so that $w_{HH} > w_{HH}^{IC}$.*

This assumption implies that compensation is shaped by agency frictions specific to continuation vehicles, allowing us to focus on these frictions without distinguishing across multiple cases. While the role of effort in the CV is crucial—shaping $w_2(z)$ and thus legacy fund compensation—time-0 effort is not critical for the legacy fund structure. We include it nonetheless for completeness, to obtain a meaningful and realistic benchmark without a continuation fund, against which we benchmark our baseline contract.

3.4 Equilibrium Summary

For clarity, we summarize the equilibrium in the following proposition:

Proposition 4 (Summary of Equilibrium)**.** *The equilibrium is characterized as follows:*

1. *At time $t = 0$, legacy LPs and the GP write a contract that stipulates payments as in Proposition 3.*

2. *At time $t = 1$, the GP exits all successful projects and raises a continuation fund only if at least one project has been successful. In that case, the GP always continues an*

M-type project, while an L-type project is continued with endogenous probability γ_z as in Proposition 2. The transfer price of asset z is given by V_z.

3. *At time $t = 2$, the GP is paid only if the CV succeeds, in which case it receives $w_2(z)$ as in Proposition 1.*

What is the source of inefficiencies in our setting? The following Proposition identifies conflicts of interest between subsets of LPs, coupled with the GP's ability to exploit them, as the driver of inefficient continuation.

Proposition 5 (Efficient Continuation)**.** *Suppose that $\xi = 1$. Thus, all legacy LPs roll over their stake and $\theta_z = 1$ upon continuation. Then, only M-type projects are continued in equilibrium, i.e., there are no inefficient zombie or crown-jewel funds.*

In Section 5, we explicitly allow for the possibility of inefficient crown-jewel funds (H-type continuation), which do not arise in the baseline. There, it becomes clear that crown-jewel funds can arise only when $\xi < 1$. Intuitively, they require benefiting at the expense of new LPs and therefore can occur only if some new LPs are present.

Overall, $\xi = 1$—that is, all legacy LPs continue and no new LPs enter—implies that continuation is NPV-efficient. In practice, however, continuing solely with legacy LPs and not bringing in new investors is likely unattractive for the GP, whose rents are reduced, and for LPs with liquidity needs or a preference for timely exit.

In Section 6.2, we show that the case $\xi = 1$ is akin to a setting where legacy fund LPs write a long-term contract with the GP at time zero.

4 Analysis and Implications

4.1 Contract Features and Conflicts of Interest

The contracts from Propositions 1 and 3 rationalize several features commonly observed in continuation and PE fund agreements (Section 1 and Kastiel and Nili, 2023), while Proposition 2 also illuminates the conflicts of interest in the context of continuation funds.

At a higher level, our analysis shows that conflicts of interest in continuation fund formation fall into two broad categories. First, conflicts between the GP and all LPs, such as the classic moral hazard regarding effort. More specific to our setting, the GP can unilaterally deviate by misrepresenting the quality of the continued firm or strategically selecting a firm

for continuation, exploiting its informational advantage over both new and legacy LPs, who cannot verify the true state.

Second, our model highlights novel conflicts of interest between subsets of LPs, where the GP aligns with one group to exploit an informational advantage over those outside the coalition. In particular, the GP may implicitly collude with (i) legacy LPs to offload a failed project onto new LPs, or (ii) new LPs to transfer a successful project at an undervalued price, thereby harming legacy LPs who exit.

We characterize the optimal contracts that address these multi-layered conflicts of interest, highlighting the features that mitigate GP rent extraction and showing how and when their absence can lead to inefficient continuations that harm subsets of LPs. In doing so, although our focus is on the optimal contract, we also illustrate how alternative contractual structures would play out.

The optimal continuation fund contract is conceptually simple and resembles real-world arrangements: the GP is rewarded upon success, resembling an equity stake with carried interest. The GP coinvests both external capital (κ) and crystallized carry from the legacy fund (explained below), which helps mitigate adverse selection but, as we will show, creates incentives for the GP to transfer successful projects at undervalued price.

The optimal legacy fund compensates the GP if at least one project is successfully exited, with higher compensation when both projects succeed. In other words, the contract links the GP's pay to both fund- and deal-level performance, a structure discussed in more detail below. Such a mix of fund- and deal-level compensation is common in practice. Our analysis suggests that this structure—whether arising from the presence of CVs or for other reasons—helps mitigate GP rent extraction associated with CVs.

No Continuation Fund Benchmark. To better understand how the option to raise a continuation vehicle shapes GP incentives, we compare the legacy fund contract from Proposition 3 to a benchmark without a continuation fund, where the model ends at time $t = 1$ and the contract is designed to minimize agency rents while incentivizing effort ($e = 1$). Appendix A.3 proves the following Proposition.

Proposition 6. *Without continuation fund, it is optimal to pay the GP only when both projects succeed, with the payment satisfying* $w_{NC} = \frac{\lambda^0}{p_A p_B}$.[23] *In addition, the following holds:*

[23] Admittedly, this optimality is not strict—alternative compensation schemes that reward the GP based on individual deal performance can yield the same payoff for LPs. However, as shown in Appendix A.3, pure fund-based compensation becomes strictly optimal if the success probability without effort is $p_z - \Delta_z$

1. *With the option to raise a continuation fund, the GP is paid for single successes, i.e.,* $w_{HL}(z), w_{HM}(z) > 0$. *When* λ^0 *is not too large (Assumption 2), these payments are higher than without a continuation fund, i.e.,* $w_{HH} > w_{NC}$.

2. *The GP's expected payoff is always higher with continuation funds, and strictly higher if* λ^0 *is not too large.*

Without a continuation fund, it is optimal to compensate the GP only if both projects succeed—that is, purely based on fund-level performance.

In contrast, under the optimal legacy fund contract in Proposition 3, the GP is compensated based on both deal-level and fund-level performance. The GP receives payment even if only one project succeeds, with compensation contingent on whether a continuation fund is raised. Pure fund-based compensation—that is, paying the GP only when both projects succeed—would encourage strategic behavior: if only one project succeeds, the GP may shift the successful project into the continuation fund while leaving underperforming ones in the legacy fund. It would also incentivize the GP to continue failed projects. Overall, to prevent the strategic transfer of both crown-jewels and zombies, the GP must be compensated based on a combination of individual deal outcomes and overall fund performance

Moreover, the option to raise a continuation fund generally leads to higher GP compensation in the legacy fund—both (1) relative to the no-continuation fund benchmark and (2) relative to the continuation fund itself. First, unlike in the no-continuation benchmark, the GP is compensated even for single-project success, and the payment for double success is higher—specifically, $w_{HH} > w_{NC}$ when λ^0 is low.

Second, compensation in the continuation fund also drives a ratcheting effect in the legacy fund. As shown in (4), (5), and (6), the GP's legacy fund compensation increases with $w_2(z)$—the amount received in the continuation fund. This higher compensation is necessary to deter the GP from transferring projects into the continuation fund solely to earn higher rewards (when the project is successful) or to capture other continuation-related benefits. If legacy fund compensation were low relative to that in the CV, the GP would transfer successful, crown-jewel projects into the continuation fund to benefit from the high compensation there.

for some $\Delta_z \in (0, p_z)$—as in Holmstrom and Tirole (1997); Maurin et al. (2023)—rather than $\Delta_z = p_z$, as assumed in our baseline. As long as $\Delta_z \in (0, p_z)$, it is strictly optimal to compensate the GP only if both projects succeed, thereby leveraging cross-pledging benefits in the spirit of Diamond (1984). These benefits vanish in the limit as $\Delta_z \to p_z$. In any case, our model predicts that while pure fund-based compensation is not optimal when a continuation fund is present, it becomes optimal (even if only weakly so) in its absence.

Overall, the GP is always better off when the option to raise a continuation fund exists. This is intuitive—continuation funds function as real options, allowing the GP to selectively transfer assets that boost their payoff. Agency rents thus increase, especially when legacy fund projects are riskier and more heterogeneous. In this sense, "the house always wins," in line with conclusions in Kastiel and Nili (2023).

Endogenous investment: Early-Crystallized Carry. The legacy fund contract compensates the GP when a single project is successfully exited, with the payment depending on whether a continuation fund is raised. The difference reflects early-crystallized carry, which arises endogenously in our model. Interestingly, $w_{HL}(z) > w_{HM}(z)$ in general, which means that the GP is paid more when no continuation fund is raised, conditional on the success of the project z. In that case, the GP receives $w_{HL}(z)$ (e.g., via carried interest) upon exit. When a continuation fund is raised, a portion of this amount—specifically, $w_{HL}(z) - w_{HM}(z)$—is crystallized early and reinvested as GP coinvestment in the continuation vehicle. Thus, the GP's coinvestment consists of both external capital (κ) and early-crystallized carry ($w_{HL}(z) - w_{HM}(z)$).

The following Corollary shows the role of early-crystallized carry and shows a substitution effect with external capital.

Corollary 1 (Early-Crystallized Carry and Endogenous Coinvestment)**.** *Suppose project $z = A$ succeeds while the other project $-z = B$ does not. Then:*

1. *Exiting both projects at $t = 1$ yields carried interest $w_{HL}(A)$, realized at the legacy fund's close if no CV is raised. If project B is transferred to a CV, then*

$$w_{HL}(A) - w_{HM}(A) = p_B w_2(B) + (p_B w_2(B) - \lambda_B)\mathbb{I}\{\gamma_B > 0\} - \kappa \geq 0$$

   dollars of carried interest are crystallized early and reinvested. The GP's total coinvestment—external funds plus early-crystallized carry—equals $w_{HL}(A) - w_{HM}(A) + \kappa \geq 0$.

2. *When $\kappa < \lambda_B$, early-crystallized carry equals $\lambda_B - \kappa$, and total coinvestment equals λ_B.*

If the GP cannot bring in sufficient external capital κ, she must reinvest a larger share of the early-crystallized carry. Thus, the lower the external capital κ, the greater the share of carry reinvested—reflecting that external capital and early-crystallized carry substitute each other so that total GP coinvestment remains unchanged. Early-crystallized carry and total coinvestment also depend on the characteristics of the continued project $-z = B$ in

the above corollary. The required coinvestment is larger, when the project faces more severe moral hazard and demands greater GP effort, such as when it is more opaque or less tangible.

4.2 GP and LP Coinvestment

A central question in the context of continuation funds is whether coinvestment by the GP or legacy LPs can mitigate agency problems.

We examine the role of LPs' and GP's coinvestment in the continuation fund. We note that the first-order determinant of LP coinvestment is the extent of liquidity shocks and capital constraints, both of which are captured by ξ. A higher ξ means that fewer LPs experience liquidity shocks, resulting in more GP coinvestment and larger θ_z, as shown in Proposition 2. The GP's coinvestment has an endogenous component via the incentive contract and an exogenous component κ, with more κ implying more funds to coinvest. The following Corollary shows how coinvestment shapes payoffs and model quantities.

Corollary 2. *The following holds:*

1. *As ξ increases, overall NPV and legacy LP payoff increase, while the GP's payoff decreases.*
2. *Suppose $\alpha = 0$. When $\kappa < \min\{\lambda_A, \lambda_B\}$, then NPV and LP payoff increase in κ, while GP payoff decreases in κ. When $\kappa > \max\{\lambda_A, \lambda_B\}$, NPV does not depend on κ and LP payoff decreases in κ. Notably, GP payoff increases in κ.*

We now provide some intuition behind these findings.

Legacy LP Coinvestment. Proposition 2 shows that θ_z—legacy LP coinvestment—increases with ξ, while γ_z—the propensity to continue failed projects—declines. Coinvestment gives legacy LPs more skin in the game, serving as a key mechanism to prevent zombie funds (i.e., the continuation of failed firms) and to mitigate adverse selection. More legacy LP coinvestment implies that fewer legacy LPs exit. This limits the GP's ability to extract rents from transferring a successful project at undervalued price, benefitting of exiting legacy LPs. As a result, greater legacy LP coinvestment lowers the GP's compensation and agency rents, improves overall NPV, and increases LP payoffs.

In this sense, coinvestment by legacy LPs helps mitigate agency conflicts in continuation vehicles—protecting new LPs from adverse selection and legacy LPs from assets being transferred at undervalued price. However, because such coinvestment reduces the GP's pay-

off, the GP has limited incentive to encourage legacy LP participation in the continuation vehicle.

GP coinvestment. Corollary 2 shows that higher κ strengthens the GP's skin-in-the-game in the continuation fund, reducing the incentive to transfer zombies. Notably, greater GP coinvestment also encourages higher legacy LP participation, as θ_z increases with κ. Intuitively, GP coinvestment signals quality, raising the continuation fund's perceived value and reducing the need to sell shares to new LPs, thereby allowing legacy LPs to retain a larger share.

However, there can be a downside of GP coinvesment: while larger GP coinvestment deters zombie transfers, it also incentivizes the GP to understate the transferred firm's quality and thus to acquire a stake in successful, continued projects at discounted prices. In practice, the GP may obtain such a larger stake either through equity coinvestment or via carried interest earnings — which are enhanced when the effective hurdle rate is lower, as a result of the firm being transferred at a lower price. To counteract this incentive, GP compensation must rise with κ. As a result, GP compensation and payoff may increase in κ, specifically, when κ is large.

In summary, GP coinvestment is a double-edged sword: on one hand, it gives the GP skin in the game, discouraging the continuation of zombie projects and mitigating adverse selection; on the other, it creates opportunities to extract rents by transferring successful projects, understating their quality, and coinvesting on favorable terms.

GP Reputation. Coinvestment provides the GP with skin in the game, aligning incentives with investors. However, it also introduces a novel agency conflict: the GP may have an incentive to opportunistically continue "crown-jewel" assets. Another important mechanism that effectively aligns incentives—by giving GPs skin in the game de facto—is reputational concern arising from repeated interactions.[24] Although we do not explicitly model reputational concerns, these are reflected in several key model parameters—for instance, λ_z (capturing the benefits of shirking, which may be tempered by reputational incentives to exert effort), or α (the GP's ability to extract surplus via side deals).

The key takeaway from our analysis is that reputational concerns—much like skin-in-the-game via coinvestment—are a double-edged sword. On the one hand, they mitigate

[24]For theoretical treatments of reputational incentives, see, e.g., Malenko and Malenko (2015); Hartman-Glaser (2017); Gryglewicz, Mayer, and Morellec (2024); for empirical evidence in the private equity context, see Ivashina and Kovner (2011); Demiroglu and James (2010).

standard agency problems (by reducing λ_z). On the other hand, a stronger reputation can enable the GP to extract greater surplus from new LPs, which may backfire by intensifying the strategic transfer of crown jewels—ultimately to the detriment of exiting legacy LPs. The GP's payoff increases unambiguously with α, while the payoff to legacy LPs declines, suggesting that greater reputational capital can benefit the GP at the expense of incumbent investors (even as new investors remain unaffected).

4.3 Project Risk and Characteristics

The following Corollary shows that the GPs' rents are highest for riskier projects, specifically, those with high upside X_z and failure risk δ_z. Loosely speaking, raising a continuation fund is an option for the GP, whose value increases with the riskiness of firms in the legacy fund.

Corollary 3. *Suppose that κ is sufficiently small. The following holds:*

1. *The GP benefits from increased failure risk, in that its payoff Π_{GP} increases with δ_z (strictly when ξ is small).*
2. *The GP benefits from increased upside, in that its payoff Π_{GP} increases in X_z.*

First, Corollary 3 shows that the GP's payoff may increase with the failure risk δ_z—even when the higher risk is not offset by a greater upside and the increase in δ_z reduces NPV. The intuition is as follows. Higher failure risk intensifies adverse selection at the continuation stage, as new LPs become more cautious about being sold zombie firms. More severe adverse selection lowers the transfer price V_z at which firm z is sold to the continuation vehicle. This lower price creates an opportunity for the GP to profit from transferring H-type firms. If the GP transfers a H-type firm while understating its quality to legacy LPs, both the GP and new LPs can acquire shares from exiting LPs at a discount. Crucially, this discount grows with the severity of adverse selection, increasing the GP's rents from transferring crown-jewels—and thus the compensation required to deter such behavior.

Second, Corollary 3 also shows that the GP's payoff increases with the potential cash flow upon success, X_z—in sharp contrast to the optimal contract without continuation funds, where the GP's payoff was independent of X_z (see Proposition 6). The underlying mechanism is similar to the above: when X_z is higher, the surplus from transferring a firm at undervalued price becomes larger. This increases the GP's ability to extract rents by opportunistically transferring successful projects at undervalued terms.

Taken together, legacy fund projects with both higher upside and higher failure risk tend to boost the GP's payoff. In other words, greater project risk benefits the GP by increasing the value of the option to raise a continuation fund—primarily through elevated agency rents. This effect operates through several channels. First, the GP's compensation exhibits convexity, much like a financial option, and therefore becomes more valuable with increased volatility. Second, the GP's discretion to selectively transfer assets into the continuation vehicle adds an additional layer of optionality. Third, higher project risk makes incentive alignment more costly, raising agency rents.

Project Heterogeneity. The possibility to raise a continuation fund features multiple layers of optionality. First, the GP has the option whether or whether not to raise a continuation fund. Second, conditional on raising a continuation fund, the GP can decide which project to transfer. The strategic decision to pick the project creates additional agency rents, making the option to raise a continuation fund more valuable when projects are ex-ante less similar. This insight is reflected the determination of w_{HH} in (6). The payment w_{HH} must be large enough to prevent continuation of either project, whichever is more favorable. Thus this payment is large if there is at least one project, for which the GP finds it favorable to continue after success. As long as the legacy fund contains such a project, incentive provision to the GP is hard and costly, while the GP can extract large agency rents.

To formally demonstrate how greater heterogeneity among legacy fund projects increases the GP's option value of raising a continuation vehicle, we consider a setting where projects are symmetric in all respects—such as success probabilities ($p_A = p_B$) and failure risks ($\delta_A = \delta_B$)—except for their cash flow potential (X_z) and private benefits (i.e., agency) parameters (λ). Specifically, we define $X_A = X + \chi\sigma$ and $X_B = X - \sigma$, while $\lambda_A = \lambda + \sigma$ and $\lambda_B = \lambda - \sigma$ where $\chi, \sigma \geq 0$ and $X > \chi\sigma$, $\lambda > \sigma$. In this formulation, project A is larger and generates higher cash flows upon success, which also require more engagement by the GP, effectively resulting in higher benefits from not exerting effort λ_A.

To clearnly illustrate that project heterogeneity boosts GP payoff, we consider the case of $\alpha = 0$ and $\kappa = 0$ to mitigate confounding effects and show:

Corollary 4. *Consider that $\kappa = \alpha = 0$. The GP's compensation w_{HH} and payoff increase in σ.*

5 Crown-Jewel Funds

In our model, the continuation of a successful project—a so-called crown-jewel fund—is inefficient. The previously optimal contract prevented such continuation, thereby maximizing NPV. However, to incentivize the GP to exit successful projects rather than transfer them to a continuation fund and coinvest on favorable terms, the contract required a relatively large reward in the legacy fund. When $\alpha = 1$, this higher reward is driven by the GP's ability to engage in a private arrangement with new LPs that rewards the GP when the continuation vehicle overperforms (i.e., produces cash flows $X_z + x > X_z$).

We now assume that rent extraction through such side deals with new LPs are not always feasible. Specifically, the parameter $\alpha \in \{0, 1\}$ in (3) is now treated as stochastic. The value of α is determined and privately observed by the GP at time $t = 1$, notably, *before* it decides whether to raise the continuation fund. Hence, α is not verifiable by LPs or contractible. It equals one with probability ϕ — where $\phi \in (0, 1)$ — and zero with probability $1 - \phi$. This captures the idea that the GP cannot always receive kickbacks from new LPs that reward continuation vehicle overperformance, or more broadly, that the GP may not always be able to strike a side arrangement—perhaps because new LPs are unwilling, unable, legally constrained, or face reputational risk.[25]

This modeling introduces an additional force. Preventing the GP from continuing H-type projects under all circumstances requires compensating the GP generously for exiting successful projects in the legacy fund—even in scenarios where it could have struck a side deal with new LPs, making continuation most tempting. As a result, the GP is effectively overpaid in cases where no such side deal is feasible. When the GP's ability to engage in side deals is not too large (i.e., ϕ is moderate), legacy LPs may occasionally prefer to allow continuation of a successful project — an outcome we denote by $\beta = 1$ (see below).

In what follows, we consider so-called "crown-jewel contracts," a second-best contract that may permit the occasional continuation of successful projects, generating "crown-jewel funds." This crown-jewel contract $\mathcal{C}^{Crown}$ pays the GP $w_{HM}(z)$, if z is exited successfully and $-z$ is continued. It pays the GP $w_{HL}(z)$, if both projects are exited but only z was successful. In this context, the payments $w_{HM}(z)$ and $w_{HL}(z)$ are those given in Proposition 3, evaluated at $\alpha = 1$. As shown there, these payments incentivize the GP exit the successful project in

[25] An alternative interpretation views α as a measure of bargaining power. In line with standard bargaining models under asymmetric information, our framework parallels the widely used random proposer model: with probability ϕ, the GP holds all bargaining power over the new LPs when proposing the side deal; with probability $1 - \phi$, the new LPs hold all the bargaining power instead.

case precisely one project has been successful, even when the GP can extract rents from a side deal with new LPs by continuing an H-type project.

The crown-jewel contract deviates from the baseline (again, considering the baseline with $\alpha = 1$) by stipulating a different payment for double success:[26]

$$w_{HH}^{Crown} := \max_{z=A,B} \{w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + (1-\beta)\Omega(z)\}. \tag{7}$$

Crucially, $\beta \in \{0, 1\}$ denotes whether the contract allows the GP to continue a successful project upon double success Thus, when $\beta = 1$, a crown-jewel fund may arise

The payment w_{HH}^{Crown} is set so that the GP finds it suboptimal to continue an H-type project when $\alpha = 0$. By contrast, if $\alpha = 1$, the GP can extract greater surplus from continuing an H-type, and continues a successful project in case both projects have been successful. Crucially, in the proof of Proposition 7, we show that the continuation stage subgame and remain identical to the baseline. Appendix A.8 presents an alternative variant of crown-jewel contracts with similar properties (but more tedious analysis), allowing the GP to continue H-type projects even in the case of single success.

The following Proposition — which assumes that λ^0 is small so that time-0 effort is not a concern (akin to assumption 2) — characterizes the crown-jewel contract, showing when allowing crown-jewel funds with $\beta = 1$ is optimal for LPs.

Proposition 7 (Crown-Jewel Funds)**.** *Under the crown-jewel contract, the following holds:*

1. *When z succeeds by time $t = 1$ and $-z$ does not, the GP exits z at time 1, as well as continues project $-z$ with probability 1 (γ_{-z}) if it is type M (type L).*
2. *When $\beta = 0$ or $\alpha = 0$, the GP always exits successful project at $t = 1$. When $\beta = 1$, the GP continues one successful project (i.e., raises a crown-jewel fund) while exiting the other if and only if α equals one. The GP generally selects for the crown-jewel fund the project z with the lower level of agency cost λ_z.*[27]
3. *$\beta = 1$ maximizes legacy LP payoff if (with z chosen for crown-jewel fund):*

$$(1-\theta_z)\left[(X_z - w_2(z))\left(1 - \frac{p_z(1-\delta_z)}{(1-\delta_z)+\delta_z\,\gamma_z}\right) + x\right] \geq \frac{\phi(I - x - \lambda_z)}{(1-\phi)}. \tag{8}$$

Interestingly, when raising a crown-jewel fund after both projects have succeeded, the GP

[26] Note that $\Omega(z)$ is from (3) and depends on $(V_z, \gamma_z, \theta_z)_{z=A,B}$ from Proposition 2.

[27] We show in the proof that project z is optimally selected whenever $\lambda_z \leq \lambda_{-z}$, except in a single knife-edge case where $\kappa > \max \lambda_A, \lambda_B$ and $\gamma_z > 0 = \gamma_{-z}$.

selects the project to transfer based on which yields the highest private payoff—reflecting the optionality embedded in continuation funds and allowing the GP to cherry-pick among projects. We show that, in general, the GP prefers the project with the lower success probability p_z and weaker agency conflicts λ_z. The intuition is that such projects enable greater rent extraction: they are sold at a lower price due to the low p_z, and limited agency frictions make surplus extraction less constrained.

In practical terms, the model predicts that when raising crown-jewel funds, the GP selects projects with a lower cost of management, as captured by a lower λ_z. A lower λ_z may also reflect reduced project opacity or risk. This aligns with the empirical observation that continuation vehicles typically involve lower management fees and carried interest for the GP compared to legacy funds.

When Do Crown-Jewel Funds Arise? Proposition 7, and in particular expression (8), identifies the conditions under which allowing the GP to raise a crown-jewel fund (i.e., $\beta = 1$) improves LP payoffs. It thus yields testable predictions about when such funds are likely to arise in practice.

Expression (8) suggests that crown-jewel funds are more beneficial to LPs—and thus more likely to arise—when θ_z is low, meaning few legacy LPs participate in the continuation fund, and when project z has high cash flow X_z but a relatively low success probability p_z or high failure probability δ_z. In other words, riskier projects with high upside and lower chances of success are more likely to be continued. The intuition is that such projects face a steeper adverse selection discount at continuation fund formation, allowing the GP to transfer crown-jewels at a lower valuation and extract greater surplus. Similarly, crown-jewel funds are more likely to emerge when the GP's stake in the continuation fund (captured by $w_2(z)$) is smaller—that is, when λ_z is low or α is high. Since both lower λ_z and higher α can reflect greater GP reputation, the model predicts that crown-jewel funds are more likely to arise when the GP is highly reputable.

Taken together, both crown-jewel funds (i.e., continuation of H-type projects) and zombie funds (continuation of L-type projects) are inefficient. They are more likely to arise when few legacy LPs roll over their stake.

Crown-Jewel Funds and Quality. When crown-jewel funds arise, the quality of transferred assets may be higher or lower than those retained in the legacy fund. As a result, the average asset quality in the continuation fund can exceed that of the legacy fund. And, when crown-jewel funds arise, they typically contain projects with lower levels of agency

cost λ_z, also a measure of quality. However, our analysis highlights an important caveat: asset quality in continuation funds is not a reliable indicator of whether LPs benefit from continuation. High asset quality may simply reflect the inefficient continuation of already successful projects, which benefits the GP and new LPs at the expense of legacy LPs.

5.1 Periods of Aggregate Liquidity Needs

What happens during periods of aggregate liquidity needs and distress? Our model provides insight into such situations, which are captured by a low value of ξ. A low ξ reflects liquidity among legacy LPs, preventing them from rolling over their commitments into the continuation fund.[28] As shown earlier, a lower fraction of rolling legacy Lps raises the GP's agency rents and increases the likelihood of continuing underperforming, zombie firms. Intuitively, when LPs face liquidity needs, they are more inclined to pressure the GP to sell assets into the continuation fund. Drawing on insights from Section 5, we also find that successful (H-type) projects are more likely to be continued when fewer legacy LPs roll.

Taken together, the model yields several predictions about CVs in times of aggregate distress. First, continuation funds are more likely to emerge when legacy LPs face high aggregate liquidity needs. While one might expect such conditions to necessarily lower the quality of firms transferred to CVs, our analysis shows otherwise. Periods of aggregate liquidity needs see heightened continuation activity driven not only by increased transfers of underperforming zombie firms, but also by the strategic retention of crown-jewel assets—alongside the continuation of M-type firms. Second, as a result, these periods tend to feature greater variation in the types of firms continued and higher variance in the quality of CVs and their portfolio companies. Third, asset quality—measured by expected cash flows—can, in fact, be higher during periods of distress.

We formalize this in the following corollary, which illustrates the mechanism analytically by comparing the limiting cases where ξ is close to one and where it is close to zero, assuming the crown-jewel contract from Proposition 7 is adopted whenever it delivers a higher payoff for legacy LPs.

Corollary 5 (Aggregate Liqudity Needs)**.** *Consider that LPs write the contract from Proposition 7 with β determined according to (8). Suppose $I + x - \lambda_z$ is not too large, and that, when an H-type project is continued, this project is z. Then, when project z is continued:*

[28] While we could explicitly model such episodes by introducing a liquidity shock realized at time $t = 1$, the qualitative implications would remain unchanged.

1. *When ξ is sufficiently large, i.e., $\xi \to 1$, project z — when continued — is of type M with probability one (i.e., $\gamma_z = 0$ and no H-type project is continued), and has expected cash flows $p_z X_z$.*

2. *When ξ is sufficiently is small, i.e., $\xi \to 0$, then $\gamma_z \geq 0$ (strictly so if κ is small), and define $q_z = 1 - (1 - p_z)\delta_z(1 - \gamma_z) - p_z(1 - \phi)$. When continued, then project z is type L with probability $\frac{\delta_z \gamma_z (1-p_z)}{q_z}$, of type M with probability $q_z^M = \frac{(1-\delta_z(1-p_z)}{q_z}$ and of type H with probability $q_z^H = \frac{p_z}{q_z}$. Expected cash flow conditional on continuation is strictly larger than $p_z X_z$ when $\phi > \gamma_z \delta_z$.*

Admittedly, in periods during which both new LPs and legacy LPs face capital constraints and liquidity needs, this result may break, since there is no buyer for continuation fund shares. Thus, the previous findings apply to situations where legacy LPs — distinct from new LPs — have high liquidity needs.

6 Additional Results

6.1 Continuation Funds of Continuation Funds

While still rare in practice, it is possible to a GP sets up a continuation fund of a continuation fund. We analyze this possibility in Appendix B.2 where we extend the model to allow the GP to continue the project further from $t = 2$ and $t = 3$. We present some analytical results in Proposition 8 below.

Crucially, we show that this possibility leads to further ratcheting up of incentives. In particular, the possibility to raise a second continuation fund implies that the GP must be paid more for success in the first continuation fund — otherwise, it continues to hold the asset (excessively long) until $t = 3$. However, the higher compensation received in the first compensation fund ratchet up compensation in the legacy fund. More formally, the possibility to raise a second continuation fund leads to higher $w_2(z)$ than in the baseline. It affects the continuation stage game (see Proposition 2) and the legacy contract (see Proposition 3) primarily through $w_2(z)$. As one can see from Proposition 3, higher $w_2(z)$ leads to higher compensation in the legacy fund — in short, the ratcheting effect is at work.

The following Proposition illustrates these results in the tractable case of symmetric projects. The proposition follows from the analysis in Appendix B.2.

Proposition 8 (Continuation Funds of Continuation Funds)**.** *Suppose projects are symmetric, κ and λ^0 are sufficiently small, and $\alpha = 0$. If the GP can raise a second continuation fund extending ownership from $t = 2$ to $t = 3$ as in Appendix B.2, then:*

1. *The GP's compensation in the first continuation fund, $w_2(z)$, is strictly higher than in the baseline.*

2. *The legacy-fund payments w_{HH}, $w_{HL}(z)$, and $w_{HM}(z)$ are also strictly higher than in the baseline.*

We also show that the possibility of raising a second continuation fund generally increases the incentive to continue zombie projects at time $t = 1$, since a failed asset can now be rolled over for two periods before liquidation. This prospect can make it more attractive for both the GP and the legacy LPs to offload failed projects onto new investors. Our analysis highlights that the option to raise a second continuation fund can directly affect the quality of assets transferred into the first.It also shows that imposing a deadline on the GP can benefit LPs by limiting rent extraction. CVs, however, weaken these positive incentive effects of deadlines—the incentives generated by deadlines for GPs are also emphasized by Zawadowski and Albuquerque (2025) in the context of private debt.

6.2 Continuation Funds via Long-Term Contracting

We contrast our setting with continuation fund with the setting in which the lifetime of the existing fund is simply extended. In our model, this scenario could be captured by focusing on long-term contracting between legacy LPs and the GP, whereby contracting determines the optimal fund duration (subject to our assumptions that only one firm can be held beyond time $t = 1$ and the fund closes if there has been no success by time 1). The following Proposition shows that long-term contracting is equivalent to our setting when $\xi = 1$ — i.e., there are no liquidity shocks.

Proposition 9 (Long-Term Contract)**.** *The optimal long-term contract is identical to our baseline solution with $\xi = 1$. It features $\gamma_z = 0$, i.e., only M-type projects are continued. The long-term contract can be implemented as the sequence of two short-term contracts, $\mathcal{C}^{Leg}$ and $\mathcal{C}^{CV}$ from Propositions 3 and 1, respectively.*

One can consider a hypothetical scenario without any liquidity shocks or capital constraints, in which the GP wishes to retain ownership of certain firms for a longer period and

faces a choice between two feasible options: (1) raising a continuation fund—where $\xi < 1$ reflects the fact that some legacy LPs are unwilling to participate even absent liquidity or capital constraints—or (2) extending the lifetime of the legacy fund, captured by $\xi = 1$.

Because the GP's payoff declines in ξ, our analysis implies that the GP strictly prefers to retain ownership through a continuation fund rather than by extending the lifetime of the existing fund. The intuition is that, in raising a continuation fund, new LPs enter, and the GP can leverage its intermediary position between exiting legacy LPs and incoming LPs to extract additional rents.

An additional consideration arises when legacy LPs' participation beyond $t = 1$ is constrained by limited capital or liquidity needs. In this case, the parameter ξ, which previously captured unwillingness to participate in continuation funds, now also reflects liquidity constraints. As a result, extending the lifetime of the legacy fund may be infeasible or costly, further reinforcing the GP's preference for the continuation fund option.

7 Final Remarks

We develop a formal theory of private equity continuation vehicles, modeling them as an embedded option available to the general partner during a fund's life. The framework integrates moral hazard and adverse selection to capture the dual role of the GP: both seller on behalf of the legacy fund and buyer on behalf of the CV. While CVs can enhance efficiency by extending the holding period of high-potential assets, they also enable inefficient continuations of failed assets and continuations of already-successful assets, reallocating value between LPs and increasing agency rents.

While our model is deliberately stylized to deliver tractable results, it necessarily abstracts from several important institutional features. One notable omission, often raised by market participants, is the use of CVs to reallocate economics within the GP organization. For example, a fund expected to deliver modest performance may generate negligible carry, particularly if carry splits favor retired or inactive partners. In such cases, both GPs and legacy LPs may be motivated by retention concerns. A CV transaction that crystallizes carry early and reallocates economics to the active team can serve to retain talent and align incentives for managing the continued asset. Our model does not explicitly capture these intra-GP dynamics, but they may be important drivers of CV formation in practice.

Likewise, we abstract from heterogeneity in LP mandates, as well as from regulatory,

tax, and jurisdictional frictions that can shape both GP and LP decisions. In reality, LPs differ in their legal capacity and willingness to roll over into a CV, and such constraints influence the investor mix, pricing, and ultimately the distribution of gains and losses. We also do not model multi-asset CVs or situations in which continuation decisions interact with concurrent fundraising for other GP vehicles—settings that could meaningfully alter the GP's incentives. A richer dynamic framework could incorporate these elements, along with repeated interactions and career concerns, to provide a more complete equilibrium characterization of GP behavior.

Even within these simplifying assumptions, the model delivers sharp, testable predictions. First, we expect higher carried interest in legacy funds relative to associated CVs, consistent with optimal ratcheting. Second, crown-jewel CVs should arise disproportionately among reputable GPs with strong bargaining power over new investors. Third, the relationship between market-wide LP liquidity shocks and CV asset quality should be non-monotonic, as such shocks can simultaneously increase the likelihood of zombie and crown-jewel transfers. Fourth, more heterogeneous legacy portfolios should produce larger GP rents and higher CV incidence. Each of these predictions can be brought to data, offering an empirical roadmap for evaluating the welfare consequences of CVs and the efficacy of governance mechanisms.

By embedding CVs in a rigorous contracting framework, we clarify the conditions under which they enhance value and when they primarily redistribute it. The core message is that CVs are neither unconditionally welfare-improving nor inherently exploitative: their efficiency hinges on contract design, the alignment of GP–LP incentives, and prevailing market conditions. The challenge is to preserve their potential for value creation while constraining their capacity for rent extraction.

References

Axelson, U., T. Jenkinson, P. Strömberg, and M. S. Weisbach (2013). Borrow cheap, buy high? the determinants of leverage and pricing in buyouts. *The journal of finance 68*(6), 2223–2267.

Axelson, U., P. Strömberg, and M. S. Weisbach (2009). Why are buyouts levered? the financial structure of private equity funds. *The Journal of finance 64*(4), 1549–1582.

Braun, R., T. Jenkinson, and C. Schemmerl (2020). Adverse selection and the performance of private equity co-investments. *Journal of Financial Economics 136*(1), 44–62.

Casamatta, C. (2003). Financing and advising: optimal financial contracts with venture capitalists. *The journal of finance 58*(5), 2059–2085.

Demiroglu, C. and C. M. James (2010). The role of private equity group reputation in lbo financing. *Journal of Financial Economics 96*(2), 306–330.

Diamond, D. W. (1984). Financial intermediation and delegated monitoring. *The review of economic studies 51*(3), 393–414.

Evercore (2025). FY 2024 secondary market review.

Gryglewicz, S. and S. Mayer (2023). Dynamic contracting with intermediation: Operational, governance, and financial engineering. *The Journal of Finance 78*(5), 2779–2836.

Gryglewicz, S., S. Mayer, and E. Morellec (2024). The dynamics of loan sales and lender incentives. *The Review of Financial Studies 37*(8), 2403–2460.

Hartman-Glaser, B. (2017). Reputation and signaling in asset sales. *Journal of Financial Economics 125*(2), 245–265.

He, Z. (2009). The sale of multiple assets with private information. *The Review of Financial Studies 22*(11), 4787–4820.

Hochberg, Y. V., A. Ljungqvist, and A. Vissing-Jørgensen (2014). Informational holdup and performance persistence in venture capital. *The Review of Financial Studies 27*(1), 102–152.

Holmstrom, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics 112*(3), 663–691.

Inderst, R., H. M. Mueller, and F. Münnich (2007). Financing a portfolio of projects. *The Review of Financial Studies 20*(4), 1289–1325.

Ivashina, V. and A. Kovner (2011). The private equity advantage: Leveraged buyout firms and relationship banking. *The Review of Financial Studies 24*(7), 2462–2498.

Ivashina, V. and J. Lerner (2019). *Patient capital: The challenges and promises of long-term investing*. Princeton University Press.

Kaplan, S. N. and P. Strömberg (2009). Leveraged buyouts and private equity. *Journal of economic perspectives 23*(1), 121–146.

Kastiel, K. and Y. Nili (2023). The rise of private equity continuation funds. *U. Pa. L. Rev. 172*, 1601.

Lazard (2025). Secondary market report 2024.

Leland, H. E. and D. H. Pyle (1977). Informational asymmetries, financial structure, and financial intermediation. *The journal of Finance 32*(2), 371–387.

Malenko, A. and N. Malenko (2015). A theory of lbo activity based on repeated debt-equity conflicts. *Journal of Financial Economics 117*(3), 607–627.

Maurin, V., D. T. Robinson, and P. Strömberg (2023). A theory of liquidity in private equity. *Management science 69*(10), 5740–5771.

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

Nadauld, T. D., B. A. Sensoy, K. Vorkink, and M. S. Weisbach (2019). The liquidity cost of private equity investments: Evidence from secondary market transactions. *Journal of Financial Economics 132*(3), 158–181.

Persaud, A. (2024). GP-led secondaries: Reshaping the landscape. Industry report, Ropes & Gray.

Schmidt, K. M. (2003). Convertible securities and venture capital finance. *The Journal of Finance 58*(3), 1139–1166.

Zawadowski, A. and R. A. Albuquerque (2025). Private credit: Risks and benefits of a maturity wall. *Available at SSRN 5189740*.

Appendix

A Proofs for Main Results

A.1 Proof of Proposition 2

We start with a more formal description of the continuation fund subgame, whereby project z is continued (while project $-z$ has succeeded). That is, we fix $z \in \{A, B\}$ — the project to be continued — in what follows. Note that z is the project index, not the type in the following signalling game.

The contract $\mathcal{C}^{Leg}$ induces truthful revelation of the state (s_A, s_B) to legacy LPs (though potential deviations by the GP must still be considered). When $s_z = H$, the GP and legacy LPs jointly maximize their payoff by exiting the project. In this case, there are no gains from trade, and continuation is inefficient—continuing an H-type project would generate a deadweight loss for the GP–legacy LP coalition. As a result, in equilibrium, only L- and M-type projects are continued.

Formally, our setting induces a signaling subgame, with two types of senders, L and M, and a receiver, the new LPs. Type M refers to the coalition of legacy LPs and the GP when the firm z they consider to transfer at time $t = 1$ is of type $s_z = M$. Analogously, type L refers to the coalition of legacy LPs and the GP when the firm they consider to transfer at $t = 1$ is of type $s_z = L$. Type M and L choose two actions, (i) whether to continue (i.e., whether to establish a continuation vehicle with firm z) or not and (ii) how much financing

$$Y_z := \max\{0, (1 - \xi)(\ell_z + I) - \kappa\}$$

to raise from new LPs, where $Y_z \geq 0$ is imposed. Recall that ℓ_z defines the endogenous time-1 payout to legacy LPs, which is a choice. Thus, choosing ℓ_z is akin to choosing Y_z (we work with the latter).

Note that the amount of financing raised, Y_z, must satisfy the constraint

$$Y_z \geq \underline{Y}_z := \max\{0, (1 - \xi)V_z - \kappa\}.$$

The lower bound $\underline{Y}_z = (1 - \xi)V_z - \kappa$ on the amount of financing raised stems from the assumption that $\ell_z \geq V_z - I$), where V_z is the endogenous share price for continuation fund equity. Note that $(1 - \xi)V_z - \kappa$ can be negative.

A.1.1 Perfect Bayesian Nash Equilibrium

We look for a Perfect Bayesian Nash equilibrium, where type M continues with probability one, while type L continues with endogenous probability $\gamma_z \in [0,1]$ — in that sense, $\gamma_z = 0$ can be viewed as separating equilibrium, while $\gamma_z \in (0,1]$ corresponds to pooling (or partial separating). When continuing and participating in the game, the L type mimics the M type by choosing the same Y_z. This occurs because the L type cannot raise financing in a separating equilibrium, since L-type projects have negative NPV.

In a Perfect Bayesian (Nash) equilibrium, each player's strategy is optimal given their beliefs (sequential rationality) — according to the payoffs introduced below — and beliefs are updated via Bayes' Rule wherever possible (belief consistency). For the two types — L and M — a strategy consists of a (binary) continuation decision (type-dependent) — that is, whether to participate in the game or not — and the choice of how much financing $Y_z \geq \underline{Y}_z$ to raise from new LPs. Y_z may in principle depend on the type, but, as we argued above, it must be in equilibrium that both types choose the same Y_z, simplifying the exposition.

Also, note that the constraint $Y_z \geq \underline{Y}_z$ restricts the action/strategy set. If the player does not continue, there is no game. Further, new LPs act passively/non-strategically in this game and break-even given their beliefs. New LPs only participate when the other player "continues" and, conditional on continuation, update their beliefs $\eta(s_z)$ after observing Y_z. Here, $\eta(s_z)$ denotes new LPs' belief that the firm to be transferred is of type s_z.

Equilibrium Conditions. Next, consider a Perfect Bayesian Nash equilibrium with $\gamma_z \in (0,1]$ and $Y_z \geq \underline{Y}_z$, where $Y_z = (1-\xi)(\ell_z + I) - \kappa \geq 0$. We denote new LPs' (equilibrium) beliefs that the project is type s_z by $\eta_z(s_z)$. Knowing that no H-type project is continued, we have $\eta(H) = 0$, by construction — indeed, the game only involves types L and M.

Anticipating the equilibrium probability γ_z with which the L-type continues/participates in the game and the prior probability δ_z for the L-type, new LPs update their beliefs according to Bayes' rule:

$$\eta_z(M) = \frac{1-\delta_z}{(1-\delta_z)+\delta_z\gamma_z} \quad \text{and} \quad \eta_z(L) = \frac{\delta_z\gamma_z}{(1-\delta_z)+\delta_z\gamma_z}.$$

We have in equilibrium that new LPs break-even given their beliefs $\eta(s_z)$, that is:

$$(1-\theta_z)V_z = Y_z, \tag{A.1}$$

where V_z satisfies

$$V_z = \underbrace{\frac{1-\delta_z}{(1-\delta_z)+\delta_z\gamma_z}}_{=\eta_z(M)} \left[p_z(X_z - w_2(z))\right]. \tag{A.2}$$

Here, $p_z(X_z - w_2(z))$ is the value of a continuation fund share if $s_z = M$, whereas the value of a continuation fund share is zero, if $s_z = L$.

When $Y_z = \underline{Y}_z$, the break-even condition (A.1) reduces to

$$(1-\theta_z)V_z = \underbrace{\max\{0, (1-\xi)V_z - \kappa\}}_{=\underline{Y}_z}.$$

Note that V_z — the fair price that new LPs pay for continuation fund shares given their beliefs — can be viewed as new LPs' best response in the signaling game.

One can solve (A.1) and (A.2) for (θ_z, V_z). Given (θ_z, V_z), the equilibrium payoff for the M type reads

$$\begin{aligned}\mathcal{M}_z &= (1-\xi)\ell_z - \xi I - \kappa + \theta_z p_z(X_z - w_2(z)) + p_z w_2(z) \\ &= Y_z - I + \theta_z \left[p_z(X_z - w_2(z))\right] + p_z w_2(z),\end{aligned} \tag{A.3}$$

where we used $Y_z = (1-\xi)(\ell_z + I) - \kappa$ to transition from the first to second line.

Analogously, the equilibrium payoff for the L type reads

$$\mathcal{L}_z = (1-\xi)\ell_z - \xi I - \kappa + \lambda_z = Y_z - I + \lambda_z. \tag{A.4}$$

The L-type chooses $\gamma_z = 1$ ($\gamma_z = 0$) and continues with probability one (probability zero) if $\mathcal{L}_z > 0$ ($\mathcal{L}_z < 0$) — as the outside option is zero. When $\mathcal{L}_z = 0$, the L-type chooses $\gamma_z \in [0,1]$. Taken together, we obtain a characterization for γ_z according to:

$$\gamma_z = \begin{cases} 1 & \text{if } \mathcal{L}_z > 0, \\ \hat{\gamma}_z \in [0,1] & \text{if } \mathcal{L}_z = 0, \\ 0 & \text{if } \mathcal{L}_z < 0. \end{cases} \tag{A.5}$$

Conditions (A.1), (A.2), and (A.5) (alongside the definition of $\mathcal{L}_z$ in (A.4)), as well as $Y_z \geq \underline{Y}_z$, are (necessary) conditions for a Perfect Bayesian Nash equilibrium. If the equilibrium exists (i.e., there exists a solution to this system with $\theta_z \in [0,1]$) the resulting Perfect Bayesian Nash equilibrium is generally not unique, as we have not pinned Y_z. We now refine the set of equilibria using the intuitive criterion, giving us a unique Perfect Bayesian Nash equilibrium under the refinement.

A.1.2 Refinement using Intuitive Criterion

Consider an equilibrium with $Y_z > \underline{Y}_z$. We show that an equilibrium with $Y_z > \underline{Y}_z$ does not survive the D1 criterion. This implies that the only possible equilibrium that survives the

D1 criterion satisfies $Y_z = \underline{Y}_z$.

Suppose that the coalition deviates by by lowering Y_z by amount ε, i.e., by setting $Y_z^\varepsilon = Y_z - \varepsilon > \underline{Y}_z$ — where $\varepsilon \in (0, Y_z - \underline{Y}_z)$. We show that, given the most favorable beliefs, the M-type is strictly better off from the deviation, whereas the L-type is strictly worse off (even under the most favorable beliefs). This implies that under the intuitive criterion refinement the unique Perfect Bayesian Nash equilibrium must feature $Y_z = \underline{Y}_z$.

In more detail, consider that after the deviation, new LPs attach (the most favorable) beliefs $\eta_z^\varepsilon(M) = 1$ to the coalition being type M; thus, new LPs believe the coalition is of type M with probability one — again, no H-type projects are continued and $\eta_z^\varepsilon(H) = 0$. Consequently, new LPs pay the fair price (which is their best response given these beliefs):

$$V_z^\varepsilon = p_z(X_z - w_2(z))$$

for continuation fund shares.

This value pins down θ_z^ε via the break-even constraint:

$$(1 - \theta_z^\varepsilon) p_z(X_z - w_2(z)) = Y_z - \varepsilon.$$

Given these beliefs, the payoff for the M type becomes (according to (A.3)):

$$\mathcal{M}_z^\varepsilon = Y_z - \varepsilon - I + \theta_z^\varepsilon \left[p_z(X_z - w_2(z))\right] + p_z w_2(z).$$

Using (A.4), the equilibrium payoff for the L type reads

$$\mathcal{L}_z^\varepsilon = Y_z - \varepsilon - I + \lambda_z.$$

Thus, the deviation gains for the M type read

$$\Delta\mathcal{M}_z^\varepsilon := \mathcal{M}_z^\varepsilon - \mathcal{M}_z = -\varepsilon + (\theta_z^\varepsilon - \theta_z) p_z(X_z - w_2(z))$$

The deviation gains for the L type read

$$\Delta\mathcal{L}_z^\varepsilon := \mathcal{L}_z^\varepsilon - \mathcal{L}_z = -\varepsilon < 0.$$

Under these favorable beliefs, namely $\eta_z^\varepsilon(M) = 1$, we have $V_z^\varepsilon = p_z(X_z - w_2(z))$ and $\theta_z^\varepsilon V_z^\varepsilon = \theta_z^\varepsilon p_z(X_z - w_2(z)) = p_z(X_z - w_2(z)) - Y_z + \varepsilon$. Therefore,

$$\Delta M_z^\varepsilon = p_z(X_z - w_2(z))(1 - \theta_z) - Y_z > (1 - \theta_z)V_z - Y_z = 0.$$

Therefore, given the most favorable beliefs attached to the deviation, the M-type strictly

prefers to deviate, while the L-type is strictly worse off deviating. Thus, any equilibrium with $Y_z > \underline{Y}_z$ does not survive the intuitive criterion.

A.1.3 Unique Perfect Bayesian Nash Equilibrium under Intuitive Criterion

This implies that the only candidate equilibrium surviving the D1 criterion satisfies $Y_z = \underline{Y}_z$, as well as (A.1), (A.2), and (A.5). We establish existence and uniqueness.

Calculate

$$\theta_z = \min\left\{\xi + \frac{\kappa}{V_z}, 1\right\} \quad \text{and} \quad V_z = \frac{(1-\delta_z)p_z}{(1-\delta_z)+\delta_z\gamma_z}\,(X_z - w_2(z)),$$

which is (2) from the main text. Clearly, V_z decreases in γ_z — this implies, all else equal, that θ_z decreases in γ_z.

From now, we focus on

$$\underline{Y}_z = (1-\xi)V_z - \kappa > 0.$$

Otherwise, when $\underline{Y}_z = 0$ and $(1-\xi)V_z - \kappa \leq 0$, the unique equilibrium features $\theta_z = 1$ and $\gamma_z = 0$.

We get $Y_z = \underline{Y}_z = (1-\xi)V_z - \kappa$ whenever $\underline{Y}_z > 0$. We can then solve

$$(1-\theta_z)V_z = (1-\xi)V_z - \kappa$$

for $\theta_z = \xi + \frac{\kappa}{V_z} \geq 0$. Moreover, we simplify:

$$\mathcal{L}_z = (1-\theta_z)V_z - I + \lambda_z = (1-\xi)V_z - \kappa - I + \lambda_z \tag{A.6}$$

One can substitute θ_z from (2) to $\mathcal{L}_z$ to express $\mathcal{L}_z$ as a function of γ_z — denoted $\mathcal{L}_z(\gamma_z)$. The equilibrium, provided it exists, is unique because $\mathcal{L}_z(\gamma_z)$ decreases in γ_z on $[0,1]$, that is, $\frac{\partial \mathcal{L}_z}{\partial \gamma_z} = (1-\xi)\frac{\partial V_z}{\partial \gamma_z} < 0$.

As for existence, we proceed constructively and solve for a candidate equilibrium. For this sake, solve

$$\mathcal{L}(\gamma_z) = (1-\xi)V_z - \kappa - I + \lambda_z = 0,$$

with $V_z = \frac{(1-\delta_z)p_z}{(1-\delta_z)+\delta_z\gamma_z}\,(X_z - w_2(z))$, for γ_z. This yields

$$\gamma_z = \hat{\gamma}_z := \frac{1-\delta_z}{\delta_z}\left[\frac{(1-\xi)p_z\,(X_z - w_2(z))}{\kappa + I - \lambda_z} - 1\right]. \tag{A.7}$$

If this expression is below zero (above one), then $\gamma_z = 0$ ($\gamma_z = 1$) in equilibrium. Otherwise, we have γ_z from (A.7), satisfying $\gamma_z \in (0,1)$.

More generally, γ_z is determined according to:

$$\gamma_z = \begin{cases} 1 & \text{if } (1-\xi)V_z - \kappa - I + \lambda_z > 0, \\ \hat{\gamma}_z \in [0,1] & \text{if } (1-\xi)V_z - \kappa - I + \lambda_z = 0, \\ 0 & \text{if } (1-\xi)V_z - \kappa - I + \lambda_z < 0, \end{cases} \tag{A.8}$$

where $\hat{\gamma}_z$ is characterized above in (A.7). Overall, a unique Perfect Bayesian Nash equilibrium surviving the intuitive criterion exists. This equilibrium is characterized by (2) and (A.8), determining $(\theta_z, \gamma_z, V_z)$.

Finally, we note that $Y_z = \underline{Y}_z$ implies that the coalition invests all funds it has available. In particular, the GP coinvests its entire time-1 endowment of κ dollars.

A.1.4 Comparative Statics

We perform comparative statics of equilibrium quantities in ξ and κ.

With respect to ξ. First, write θ_z and γ_z as functions of ξ, i.e., $\theta_z(\xi), \gamma_z(\xi)$. Note that V_z depends on ξ only via $\gamma_z(\xi)$.

Assume that $\gamma_z(\xi)$ and $\theta_z(\xi)$ are differentiable — which is the case almost everywhere. Next, calculate using (A.6):

$$\frac{d\mathcal{L}_z}{d\xi} = -\theta_z'(\xi)V_z + (1-\theta_z)\frac{\partial V_z}{\partial \gamma_z}\gamma_z'(\xi). \tag{A.9}$$

Note that $\frac{\partial V_z}{\partial \gamma_z} < 0$. In addition, we use $\theta_z = \xi + \frac{\kappa}{V_z}$ to calculate:

$$\theta_z'(\xi) = 1 - \frac{\kappa}{(V_z)^2}\frac{\partial V_z}{\partial \gamma_z}\gamma_z'(\xi). \tag{A.10}$$

If $\gamma_z \in \{0,1\}$, then $\gamma_z'(\xi) = 0$ and $\theta_z'(\xi) > 0$, so $\frac{d\mathcal{L}_z}{d\xi} < 0$.

If $\gamma_z \in (0,1)$, then the we have $\mathcal{L}_z \equiv 0$ and $\frac{d\mathcal{L}_z}{d\xi} = 0$. By (A.9), we then get $\theta_z'(\xi)V_z = (1-\theta_z)\frac{\partial V_z}{\partial \gamma_z}\gamma_z'(\xi)$. Thus, owing to $\frac{\partial V_z}{\partial \gamma_z} < 0$, it follows that $\theta_z'(\xi)$ and $\gamma_z'(\xi)$ have opposite signs. If it were $\gamma_z'(\xi) > 0$, then $\theta_z'(\xi) < 0$, a contradiction to (A.10). Thus, it must be that $\theta_z'(\xi) > 0$ and $\gamma_z'(\xi) < 0$ when $\gamma_z \in (0,1)$.

Combined with $\theta_z'(\xi) > 0$ as well as $\frac{d\mathcal{L}_z}{d\xi} < 0$ for $\gamma_z \in \{0,1\}$ (note that $\mathcal{L}_z$ determines γ_z according to (A.8)), it follows that θ_z increases in ξ while γ_z decreases in ξ.

With respect to κ. Second, analogous to before, write θ_z and γ_z as functions of κ, i.e., $\theta_z(\kappa)$ and $\gamma_z(\kappa)$. Assume $\gamma_z(\kappa)$ and $\theta_z(\kappa)$ are differentiable, so that:

$$\frac{d\mathcal{L}_z}{d\kappa} = (1-\xi)\frac{\partial V_z}{\partial \gamma_z}\gamma_z'(\kappa) - 1. \tag{A.11}$$

As before, $V_z > 0$ and $\frac{\partial V_z}{\partial \gamma_z} < 0$.

If $\gamma_z \in \{0,1\}$, then $\gamma_z'(\kappa) = 0$ and $\theta_z'(\kappa) > 0$, as well as $\frac{d\mathcal{L}_z}{d\kappa} < 0$. Otherwise, we have $\frac{d\mathcal{L}_z}{d\kappa} = 0$ and therefore — by means of (A.11) — $\gamma_z'(\kappa) < 0$. Overall, γ_z decreases in κ.

Next, we use $\theta_z = \xi + \frac{\kappa}{V_z}$ to calculate:

$$\theta_z'(\kappa) = \frac{1}{V_z}\left(1 - \frac{\kappa}{V_z}\frac{\partial V_z}{\partial \gamma_z}\gamma_z'(\kappa)\right). \tag{A.12}$$

This readily implies $\theta_z'(\kappa) > 0$.

Combining all cases, θ_z increases in κ, while γ_z decreases in κ.

A.2 Proof of Proposition 3

We proceed in several steps. First, we characterize the incentive constraints for the legacy fund contract $\mathcal{C}^{Leg}$, ensuring truthful reporting, full effort, and a continuation strategy chosen by the GP consistent with the continuation stage subgame.

A.2.1 Incentive Constraints

We now derive all the incentive constraints for the legacy fund contract $\mathcal{C}^{Leg}$. The contract is incentive-compatible if it induces (i) full effort ($e = 1$), (ii) truthful revelation of the state s_z to legacy LPs (i.e., $\hat{s}_z = s_z$ for $z = A, B$), and (iii) a continuation strategy consistent with the continuation stage game. Regarding (iii), conditional on z succeeding, the contract must incentivize the GP to exit project z at time $t = 1$, and:

1. Exit project $-z$ at time $t = 1$ if $s_{-z} = H$
2. Continue project $-z$ with probability one if $s_{-z} = M$
3. Continue project $-z$ with probability γ_{-z}, as determined in the continuation stage game, if $s_{-z} = L$

To derive these incentive constraints, we take the continuation game quantities

$$(V_z, \gamma_z, w_2(z))_{z=A,B}$$

as given. These are determined in the continuation stage subgame (see Proposition 2), with the legacy fund contract inducing the GP implement a consistent continuation strategy.

We distinguish between several cases.

Case 1: One project succeeded, the other one is high-potential ($s_z = H, s_{-z} = M$). Suppose project z has succeeded, i.e., $s_z = H$, while $s_{-z} = M$, so that the contract prescribes that z is exited and $-z$ is continued. This yields expected payoff for the GP of $w_{HM}(z)$ (from exiting z) and $p_{-z}w_2(-z)$ dollars in expectation from the continuation fund while working, while the GP coinvests κ external funds.

A potential deviation is to continue z instead of $-z$. This deviation is most tempting when new investors are informed. In this case, when continuing z, the GP can collude with new LPs and extract surplus $\alpha\Omega(z)$. In addition, the GP receives payment $w_2(z)$ with certainty at $t = 2$ and can shirk when operating the continuation fund, deriving private benefits of λ_z.

In combination, we obtain that the GP does not deviate by continuing z instead of $-z$ if and only if

$$w_{HM}(z) \geq w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z}w_2(-z), \tag{IC-1A}$$

which must hold for both $z = A, B$ for $\mathcal{C}^{Leg} \in IC$.

In addition, the GP must be better off not stopping, in that

$$w_{HM}(z) + p_{-z}w_2(-z) - \kappa \geq w_{HL}(z). \tag{IC1-B}$$

Case 2: One project succeeded, the other one failed ($s_z = H, s_{-z} = L$). Next, suppose project z has succeeded, while $-z$ has failed. When $\gamma_{-z} = 0$, i.e., the contract prescribes to exit $-z$, it is necessary to incentivize the GP to disclose failure of $-z$, yielding a payment of $w_{HL}(z)$ from successful exit of project z. If the GP pretends $-z$ is of type M and raises a continuation fund, it co-invests κ dollars, receives a payment of $w_{HM}(z)$ from exiting z. In the continuation fund, the GP then shirks on $-z$, deriving private benefits λ_{-z}.

Taken together, the GP is better off disclosing failure (as opposed to hiding it) if:

$$w_{HL}(z) \geq w_{HM}(z) + \lambda_{-z} - \kappa.$$

When $\gamma_{-z} \in (0, 1)$, the GP must be indifferent between disclosing failure or not this must hold in equality — which will be the case. When $\gamma_{-z} = 1$, above inequality is reversed. More generally, the incentive constraint regarding γ_{-z} — the disclosure and continuation of

a failed project — becomes

$$w_{HL}(z)\begin{cases}\geq w_{HM}(z)+\lambda_{-z}-\kappa, & \text{if } \gamma_{-z}=0\\ =w_{HM}(z)+\lambda_{-z}-\kappa, & \text{if } \gamma_{-z}\in(0,1)\\ \leq w_{HM}(z)+\lambda_{-z}-\kappa, & \text{if } \gamma_{-z}=1.\end{cases} \tag{IC-2}$$

A similar incentive problem and incentive constraint arises in Mayer (2022), essentially requiring that the GP must be rewarded to disclose failure.

Suppose again that project z has succeeded, while $-z$ failed. Another potential deviation is to continue z, which yields, in the best case when new LPs are informed the high quality of z, payoff of $w_2(z)+\alpha\Omega(z)+\lambda_z-\kappa$. Whenever the contract, however, stipulates to exit both z and $-z$ at time $t=1$ with some positive probability (i.e., when $\gamma_{-z}<1$), then the GP must better off from stopping altogether (and exiting z successfully) and earning $w_{HL}(z)$ than from continuing z. This leads to the incentive constraint

$$w_{HL}(z)\geq w_2(z)+\alpha\Omega(z)+\lambda_z-\kappa \quad \text{for} \quad \gamma_{-z}<1. \tag{IC-3A}$$

We note that when $\gamma_z\in(0,1)$, the GP is indifferent between continuing $-z$ (while exiting z) and exiting altogether.

At the same time, the GP must find it worthwhile to continue $-z$ rather than z when the contract prescribes to do so with some probability, i.e., when $\gamma_{-z}>0$. This leads to the incentive constraint:

$$w_{HM}(z)+\lambda_{-z}\geq w_2(z)+\alpha\Omega(z)+\lambda_z \quad \text{for} \quad \gamma_{-z}>0. \tag{IC-3B}$$

Case 3: Both project succeeded ($s_z=H, s_{-z}=H$). Suppose that both projects have succeeded. It is then prescribed to exit both projects, which generates payout w_{HH} to the GP. However, the GP could deviate and continue any project z, which is most tempting when new LPs are informed. Continuing project z then allows the GP to extract surplus $\alpha\Omega(z)$ from colluding with new LPs, in addition to receiving payment $w_2(z)$ at time $t=2$. Further, the GP derives benefits of λ_z, but needs to co-invest κ dollars. Thus, the GP is better off exiting both projects if and only if for $z=A,B$:

$$w_{HH}\geq w_{HM}(-z)-\kappa+\lambda_z+w_2(z)+\alpha\Omega(z). \tag{IC-4}$$

Note that in this deviation, the GP can pick which of the two successful projects to continue.

Case 4: Time-0 Effort Incentives. Define $p_z(e) = p_z - (1-e)p_z$ for $z = A, B$ the success probability depending on effort at time $t = 0$. Anticipating the continuation stage and continuation fund contracting, the GP's initial payoff under contract $\mathcal{C}^{Leg}$ and efforts e reads:

$$\begin{aligned}\Pi_{GP}(e;\mathcal{C}^{Leg}) = p_A(e)p_B(e)w_{HH} + \lambda^0(1-e) \\ + \sum_{z=A,B} p_{-z}(e)(1-p_z(e))\Big\{(1-\delta_z)\big[p_z w_2(z) - \kappa + w_{HM}(-z)\big] \\ + \delta_z\big[(1-\gamma_z)w_{HL}(-z) + \gamma_z\big(-\kappa + w_{HM}(-z) + \lambda_z\big)\big]\Big\}.\end{aligned} \tag{A.13}$$

The effort incentive constraint then reads

$$1 \in \arg\max_{e\in\{0,1\}} \Pi_{GP}(e;\mathcal{C}^{Leg}). \tag{IC-5}$$

This incentive constraint can therefore be rewritten as $w_{HH} \geq w_{HH}^{IC}(\mathcal{C}^{Leg})$ where

$$w_{HH}^{IC}(\mathcal{C}^{Leg}) = \frac{\max\{0, \mathcal{W}_{HH}^{IC}\}}{p_A p_B}. \tag{A.14}$$

for

$$\begin{aligned}\mathcal{W}_{HH}^{IC}(\mathcal{C}^{Leg}) := -\sum_{z=A,B} p_{-z}(1-p_z)\Big[(1-\delta_z)[p_z w_2(z) - \kappa + w_{HM}(-z)] \\ + \delta_z\Big(\gamma_z\{w_{HM}(-z) + \lambda_z - \kappa\} + (1-\gamma_z)w_{HL}(-z)\Big)\Big] + \lambda^0.\end{aligned}$$

Note that $w_{HH}^{IC}(\mathcal{C}^{Leg})$ depends on other contract terms, with $w_{HH}^{IC}(\mathcal{C}^{Leg}) \leq \frac{\lambda^0}{p_A p_B}$, as long as the other incentive constraints are met. Then, w_{HH}^{IC} from the statement of the Proposition is the defined by inserting the payments $w_{HM}(z)$ and $w_{HL}(z)$ into (A.14).

Summary. We conclude that the contract $\mathcal{C}^{Leg}$ is incentive compatible — that is, $\mathcal{C}^{Leg} \in IC$ — if and only if the incentive conditions (IC-1A), (IC1-B), (IC-2), (IC-3A), (IC-3B), (IC-4), and (IC-5) are satisfied.

A.2.2 Optimal Legacy Fund Contract

Given the structure outlined before, the contract $\mathcal{C}^{Leg}$ of the legacy fund can be represented as $\mathcal{C}^{Leg} = (w_{HH}, w_{HM}(z), w_{HL}(z))_{z=A,B}$, which induces the prescribed actions (e.g., γ_z) and truthful disclosure of the state to legacy investors.

Payoffs and Optimization. Under incentive compatible contracts, the NPV reads

$$NPV(\mathcal{C}^{Leg}) = p_A p_B \left(X_A + X_B\right) + \sum_{z=A,B} p_{-z}(1-p_z)\Big\{X_{-z} + (1-\delta_z)(p_z X_z - I) + \delta_z \gamma_z \big[\lambda_z - I\big]\Big\}, \tag{A.15}$$

where γ_z is from (A.8).

Notably, the NPV is independent of the exact contract $\mathcal{C}^{Leg}$ as long as $\mathcal{C}^{Leg} \in IC$. When writing down existing LPs payoff, we note that all players are risk-neutral without discounting and utilize that new LPs make zero expected profits in equilibrium (i.e., they break even). Thus, LPs' payoff becomes:

$$\Pi_{LP} = \max_{\mathcal{C}^{Leg} \in IC} \big[NPV(\mathcal{C}^{Leg}) - \Pi_{GP}(1; \mathcal{C}^{Leg})\big], \tag{A.16}$$

where $\Pi_{GP}(1; \mathcal{C}^{Leg})$ was defined in (A.13).

Thus, to optimize legacy LPs' payoff, we minimize the GP's rent subject to all incentive conditions (IC-1A)-(IC-5). To minimize agency rents, it is optimal to minimize payments to the GP subject to all relevant incentive constraints — a linear programming problem. In this process, we are taking z (the project potentially being continued) and γ_z as given.

The incentive constraints admit a hierarchical interdependence, whereby (IC-1A) only depends on the legacy contract only via $w_C(z)$, (IC-2) depends on the legacy contract only via $w_C(z), w_{HL}(z)$, and (IC-3A) and (IC-4) depend on the legacy contract via $w_C(z), w_{HL}(z), w_{HH}$.

Combining (IC-1A) and (IC-3B). We start by combining (IC-1A) and (IC-3B) in order to solve for the minimum payment $w_{HM}(z)$ satisfying both constraints:

$$w_{HM}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z} w_2(-z) + (p_{-z} w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}, \tag{A.17}$$

where $p_{-z} w_2(-z) \geq \lambda_{-z}$ by (IC-C). Thus, $(p_{-z} w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\} \geq 0$.

Combining (IC-2) and (IC-3A). Next, given $w_C(z)$ for $z = A, B$ determined above, we turn to incentive constraints (IC-2) and (IC-3A) to solve for $w_{HL}(z)$.

First, consider $\gamma_{-z} = 1$. Then, $w_{HL}(z)$ can be set arbitrarily (e.g., to zero) and is payoff-irrelevant, as long as it induces the GP to always continue (i.e., (IC1-B) and (IC-3B) are met).

Second, consider $\gamma_{-z} < 1$. Under these circumstances, the GP must find it (weakly) optimal to stop at time $t = 1$ and to earn payoff $w_{HL}(z)$, in case z has succeeded and $-z$ has failed — that is, (IC-3A) and (IC-2) must hold. Next solve for the minimum payment

$w_{HL}(z)$ satisfying (IC-3A):

$$w_{HL}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa \tag{A.18}$$

Calculate using the expression for $w_{HM}(z)$ from (A.17):

$$w_{HL}(z) = w_{HM}(z) + p_{-z}w_2(-z) - (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\} - \kappa, \tag{A.19}$$

which will be used below. We conjecture this payment $w_{HL}(z)$ and verify next that this choice of $w_{HL}(z)$ implies that incentive constraints (IC-2) and (IC1-B) are being satisfied.

Verifying (IC-2) and (IC1-B). We start by verifying that (IC-2) is met under $w_{HL}(z)$ from (A.18). First, when $\gamma_{-z} = 0$, then we get using (A.19)

$$w_{HL}(z) = w_{HM}(z) + p_{-z}w_2(-z) - \kappa \geq w_{HM}(z) + \lambda_{-z} - \kappa$$

where we used that $\lambda_{-z} \leq p_{-z}w_2(-z)$ (see (IC-C) and (PC-C) in Proposition 1). Thus, (IC-2) is satisfied under $w_{HL}(z)$ from (A.18).

Second, when $\gamma_{-z} \in (0, 1)$, then we get using (A.19)

$$w_{HL}(z) = w_{HM}(z) + \lambda_{-z} - \kappa$$

and (IC-2) holds in equality. Third, it can be see that for $\gamma_{-z} = 1$, incentive constraint (IC-2) holds too.

Finally, we verify that under our proposed payments, the constraint (IC1-B) is met. For this sake, use (A.19) to obtain

$$w_{HL}(z) = w_{HM}(z) + p_{-z}w_2(-z) - (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\} - \kappa \leq w_{HM}(z) + p_{-z}w_2(-z) - \kappa$$

which is incentive constraint (IC1-B).

Combining (IC-4) and (IC-5). Finally, given these payments $w_{HL}(z)$ and $w_{HM}(z)$, we can calculate $w_{HH}^{IC}(\mathcal{C}^{Leg})$ according to (A.14) — let w_{HH}^{IC} the value of $w_{HH}^{IC}(\mathcal{C}^{Leg})$ under $w_{HM}(z)$ from (A.17) and $w_{HL}(z)$ from (A.18). Thus, incentive constraint (IC-5) can be written as $w_{HH} \geq w_{HH}^{IC}$.

Note that it is optimal to choose for w_{HH} the minimum value satisfying $w_{HH} \geq w_{HH}^{IC}$ — that is, (IC-5) and (IC-4) — which leads to

$$w_{HH} = \max_{z=A,B} \left\{ w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z), w_{HH}^{IC} \right\}.$$

The argument is hereby complete and the optimal contractual payments are therefore as stated in Proposition 3.

Simplification. Using (A.17), we rewrite the first term in square brackets as follows:

$$\begin{aligned} w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z) &= w_2(-z) + \alpha\Omega(-z) + \lambda_{-z} - p_z w_2(z) + \alpha\Omega(z) \\ &\quad + (p_z w_2(z) - \lambda_z)\mathbb{I}\{\gamma_z > 0\} + \lambda_z + w_2(z) \\ &= \big[w_2(A) + w_2(B) + \lambda_A + \lambda_B + \alpha\Omega(A) + \alpha\Omega(B)\big] - p_z w_2(z) + (p_z w_2(z) - \lambda_z)\mathbb{I}\{\gamma_z > 0\} \end{aligned}$$

Thus:

$$\begin{aligned} &\max_{z=A,B} \{w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z)\} \\ &= \big[w_2(A) + w_2(B) - \kappa + \lambda_A + \lambda_B + \alpha\Omega(A) + \alpha\Omega(B)\big] \\ &\qquad - \min_{z=A,B} \Big[p_z w_2(z) - (p_z w_2(z) - \lambda_z)\mathbb{I}\{\gamma_z > 0\}\Big]. \end{aligned} \tag{A.20}$$

We will use this relationship later on.

A.2.3 Payoffs

Under the optimal contract, the GP's payoff can be written as:

$$\begin{aligned} \Pi_{GP} &:= p_A p_B w_{HH} \\ &\quad + \sum_{z=A,B} p_{-z}(1 - p_z)\Big\{(1 - \delta_z)[p_z w_2(z) - \kappa + w_{HM}(-z)] + \delta_z w_{HL}(-z)\Big\}. \end{aligned} \tag{A.21}$$

Suppose z succeeded, while $-z$ failed. One can show that, under the payment scheme stipulated by the optimal contract, the GP strictly prefers to terminate the failed project $-z$ when $\gamma_{-z} = 0$, and is indifferent between continuation and termination when $\gamma_{-z} \in (0, 1]$. Indeed, when $\gamma_{-z} > 0$, continuing project $-z$ yields

$$w_{HM}(z) + \lambda_{-z} - \kappa = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa = w_{HL}(z), \tag{A.22}$$

after using that $w_{HM}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z} w_2(-z) + (p_{-z} w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}$. Accordingly, for notational convenience we shall henceforth present the GP's payoff "as if" she always terminates the project upon failure.

The NPV under the optimal contract is given by:

$$NPV = p_A p_B \left(X_A + X_B\right) + \sum_{z=A,B} p_{-z}(1-p_z)\Big\{X_{-z} + (1-\delta_z)(p_z X_z - I) - \delta_z \gamma_z \big[I - \lambda_z\big]\Big\}. \tag{A.23}$$

As such, legacy LPs' payoff is

$$\Pi_{LP} = NPV - \Pi_{GP}, \tag{A.24}$$

while new LPs break-even in expectation.

A.3 No Continuation Fund Benchmark: Proof of Proposition 6

We present the solution to the no continuation fund benchmark. We assume that not exerting effort reduces the probability of success from p_z (with effort) to $p_z - \Delta_z$ for $\Delta_z \in (0, p_z)$. We show that it is strictly optimal, for any $\Delta_z \in (0, p_z)$, that the GP is only paid upon double success. We then take the limit $\Delta_z \to p_z$. As is well known, when $\Delta_z = p_z$, the strict optimality vanishes and it is weakly optimal to pay the GP upon double success. However, we focus on the contract that survives the limit argument, which uniquely pays the GP only upon double success.

A.3.1 Solution and Optimal Contract

The LPs design a contract $\mathcal{C}^{NC} = (w_z, w_{HH}^{NC})_{z=A,B}$ — which pays the GP $w_z \geq 0$ in case z succeeds and w_{HH}^{NC} in case both projects succeed. Let $p_z(e) = p_z - \Delta_z(1-e)$.

The GP's payoff given this contract and effort choice $e \in \{0,1\}$ equals:

$$\Pi_{GP}^{NC}(e, \mathcal{C}^{NC}) = p_A(e)(1-p_B(e))w_A + p_B(e)(1-p_A(e))w_B + p_A(e)p_B(e)w_{HH}^N C + \lambda^0(1-e).$$

The effort incentive constraint $1 \in \arg\max_{e\in\{0,1\}} \Pi_{GP}^{NC}(e, \mathcal{C}^{NC})$ can therefore be rewritten as:

$$\begin{aligned} w_{HH}^{NC} p_A p_B + w_A p_A(1-p_B) + w_B p_B(1-p_A) \geq w_{HH}^{NC}(p_A - \Delta_A)(p_B - \Delta_B) \\ + w_A(p_A - \Delta_A)(1 - p_B + \Delta_B) + w_B(p_B - \Delta_B)(1 - p_A + \Delta_A) + \lambda^0. \end{aligned}$$

Thus:

$$\begin{aligned} w_{HH}^{NC}\big(p_A\Delta_B + p_B\Delta_A - \Delta_A\Delta_B\big) + w_A\big[\Delta_A\left(1 - p_B + \Delta_B\right) \;-\; p_A\,\Delta_B\big] \\ + w_B\big[\Delta_B\left(1 - p_A + \Delta_A\right) \;-\; p_B\,\Delta_A\big] \;\geq\; \lambda^0. \end{aligned} \tag{A.25}$$

The LP's payoff reads then:

$$\Pi_{LP}^{NC} = (p_A X_A + p_B X_B) - p_A p_B w_{HH}^{NC} - p_A(1-p_B)w_A - p_B(1-p_A)w_B. \tag{A.26}$$

This payoff is maximized over $\mathcal{C}^{NC}$ subject to (A.25). There is no relevant participation constraint, since payments are positive and the GP has an outside option of zero.

The IC constraint optimally binds, so that

$$w_{HH}^{NC} \;=\; \frac{\lambda^0 - w_A\big[\Delta_A\,(1-p_B+\Delta_B) \;-\; p_A\,\Delta_B\big] - w_B\big[\Delta_B\,(1-p_A+\Delta_A) \;-\; p_B\,\Delta_A\big]}{p_A\,\Delta_B + p_B\,\Delta_A - \Delta_A\,\Delta_B}\,.$$

Inserting this expression for w_{HH}^{NC} into (A.26) and simplifying, we obtain:

$$\begin{aligned}\Pi_{LP}^{NC} = p_A X_A + p_B X_B - \frac{p_A p_B\,\lambda^0}{p_A\Delta_B + p_B\Delta_A - \Delta_A\Delta_B} \\ + \sum_{z=A,B} w_z\Big(\frac{p_z p_{-z}\big[\Delta_z(1-p_z+\Delta_{-z}) - p_z\Delta_{-z}\big]}{p_z\Delta_{-z} + p_{-z}\Delta_z - \Delta_z\Delta_{-z}} - p_z(1-p_z)\Big).\end{aligned} \tag{A.27}$$

Next, we show that:

$$\mathcal{D} := \Big(\frac{p_z p_{-z}\big[\Delta_z(1-p_z+\Delta_{-z}) - p_z\Delta_{-z}\big]}{p_z\Delta_{-z} + p_{-z}\Delta_z - \Delta_z\Delta_{-z}} - p_z(1-p_z)\Big) \le 0.$$

This follows from

$$p_z p_{-z}\big[\Delta_z(1-p_{-z}+\Delta_{-z}) - p_z\Delta_{-z}\big] - p_z(1-p_{-z})\big[p_z\Delta_{-z} + p_{-z}\Delta_z - \Delta_z\Delta_{-z}\big] = -p_z\Delta_{-z}(p_z-\Delta_z) \le 0,$$

where we used $p_z \ge \Delta_z$. The inequality is strict for $\Delta_z < p_z$. Thus, the objective in (A.27) is maximized upon setting $w_z = 0$, i.e., the GP is only paid for double success, where

$$w_{HH}^{NC} \;=\; \frac{\lambda^0}{p_A\,\Delta_B + p_B\,\Delta_A - \Delta_A\,\Delta_B}\,.$$

The GP's payoff then becomes $p_A p_B w_{HH}^{NC}$.

In the limit $\Delta_z \to p_z$, we then get $\lim_{\Delta_z \to p_z} w_{HH}^{NC} = \frac{\lambda^0}{p_A p_B}$, as desired.

A.3.2 GP Always Earns Higher Payoff

The GP's payoff absent continuation fund equals (for $\Delta_z = p_z$): $\Pi_{GP}^{NC} = \lambda^0$.

Under the optimal contract $\mathcal{C}^{Leg}$ (with continuation fund), the GP's payoff satisfies (A.21). Note that $w_{HH} \ge w_{HH}^{IC}$ by incentive constraint (IC-5) — where $w_{HH}^{IC} = w_{HH}^{IC}(\mathcal{C}^{Leg})$ is defined in (A.14). Thus, the GP's payoff with continuation fund, $\Pi_{GP} = \Pi_{GP}(1; \mathcal{C}^{Leg}$, is

bounded from below by

$$\underline{\Pi} := p_A p_B w_{HH}^{IC} + \sum_{z=A,B} p_{-z}(1-p_z)\Big[(1-\delta_z)[p_z w_2(z) - \kappa + w_{HM}(-z)] + \delta_z w_{HL}(-z)\Big].$$

where the inequality is strict if $w_{HH} > w_{HH}^{IC}$ — which is always the case for small λ^0.

By (IC-5), we necessarily have $\underline{\Pi} \geq \lambda^0$ Therefore:

$$\Pi_{GP} \geq \Pi_{GP}^{NC},$$

where the inequality is strict if $w_{HH} > w_{HH}^{IC}$, which is the case when λ^0 is not too large.

A.4 Proof of Corollary 2

For the comparative statics, we assume that λ^0 is such that constraint (IC-5) does not bind.

The payments read:

$$\begin{aligned}
w_{HM}(z) &= w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\},\\
w_{HL}(z) &= w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa,\\
w_{HH} &= \max_{z=A,B}\left\{w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z)\right\}.
\end{aligned}$$

Under the optimal contract, the GP's payoff Π_{GP} satisfies (A.21), the LP payoff satisfies (A.24), while NPV satisfies (A.23)

A.4.1 Comparative Statics with respect to ξ

Observe that Π_{GP} depends on ξ only via $\Omega(z)$ and γ_z — for $z = A, B$. Note that $\Omega(z)$ is defined in (3) — it equals zero when $\theta_z = 1$, i.e., when $\kappa \geq (1-\xi)V_z$.

First, note that γ_z matters for the payoff via $w_{HM}(-z)$. Since $p_z w_2(z) \geq \lambda_z$ by Proposition 1, we have that an increase in γ_z — all else equal — increases $w_{HM}(-z)$. Since γ_z declines with ξ (see Proposition 2), the effect of ξ — operating through γ_z — is (weakly) negative.

Second, consider how ξ affects Π_{GP} via $\Omega(z)$ for $z = A, B$. Consider $\alpha = 1$ and $1-\theta_z > 0$ for at least one z — otherwise, $\alpha\Omega(z) = 0$ for both $z = A, B$ and Π_{GP} is not sensitive to ξ. Also, note that, holding all else equal, $w_{HL}(z)$, $w_{HM}(z)$, and w_{HH} increase in $\Omega(z)$. Thus, $\frac{\partial \Pi_{GP}}{\partial \xi}$ has the same sign as $\frac{\partial \Omega(z)}{\partial \xi}$. We calculate:

$$\frac{d\Omega(z)}{d\xi} = -(1-\theta_z)\frac{dV_z}{d\xi} - (X_z + x - w_2(z) - V_z)\frac{d\theta_z}{d\xi}.$$

where $X_z + x - w_2(z) - V_z > 0$. By Proposition 2, we have $\frac{\partial \theta_z}{\partial \xi} > 0$. Further, note that V_z depends on ξ only via γ_z so that (provided differentiability):

$$\frac{dV_z}{d\xi} = \frac{\partial V_z}{\partial \gamma_z} \frac{d\gamma_z}{d\xi}$$

Proposition 2 implies $\frac{d\gamma_z}{d\xi} \leq 0$ (whenever differentiable), which, combined with $\frac{\partial V_z}{\partial \gamma_z} < 0$, yields $\frac{dV_z}{d\xi}$.

Overall, we obtain that $\Omega(z)$ and the GP's payoff decrease in ξ.

Next, observe that the NPV increases in ξ, since γ_z declines in ξ. This means that LP payoff increases in ξ too.

A.4.2 Comparative Statics with respect to κ

Note that κ affects Π_{GP} in three ways: (i) (directly) through the payments $w_{HM}(z), w_{HL}(z), w_{HH}$, (ii) through the payments via $\Omega(z)$, and (iii) directly and negatively in expression (A.21). To begin with note that regarding (iii), the direct effect of κ on Π_{GP} is negative, since Π_{GP} decreases in κ holding contract payments fixed. Moreover, since $\alpha = 0$, the effect (ii) through $\alpha\Omega(z)$ is zero too.

Small κ. First, suppose that κ is small, i.e., $\kappa < \min\{\lambda_A, \lambda_B\}$. Then, $p_z w_2(z) = \lambda_z$ for $z = A, B$. Note that a marginal increase in κ reduces $w_{HL}(z)$ and w_{HH}, holding all else (including $\Omega(z)$) equal. Thus, Π_{GP} declines in κ.

The NPV increases in κ. As a result, LP payoff increases in κ too.

Large κ. Second, suppose that κ is large, so that $\kappa > \max\{\lambda_A, \lambda_B\}$. Due to $\alpha\Omega(z) = 0$, payments simplify to

$$w_{HM}(z) = \frac{\kappa(1-p_z)}{p_z} + \lambda_z; \quad w_{HL}(z) = \frac{\kappa(1-p_z)}{p_z} + \lambda_z$$
$$w_{HH} = \max_{z=A,B} \left\{ w_{HM}(-z) - \kappa + \lambda_z + \frac{\kappa}{p_z} \right\}$$

Note that $w_{HM}(z)$ and $w_{HL}(z)$ strictly increase in κ, and therefore also w_{HH} increases in κ. Since $p_z w_2(z) - \kappa = 0$, the GP's payoff in (A.21) depends on κ only via $w_{HM}(z)$, $w_{HL}(z)$, and w_{HH}. It increases in $w_{HM}(z)$, $w_{HL}(z)$, and w_{HH}. Thus, Π_{GP} increases in κ.

Under these circumstances, we have $\gamma_z = 0$ and thus NPV no more depends on κ. This implies that LP payoff decreases in κ.

A.5 Proof of Corollary 3

Consider that $\kappa = 0$. By continuity, the statements hold for sufficiently small $\kappa \geq 0$. Note $p_z w_2(z) = \lambda_z$. This relationship implies

$$w_{HL}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa = w_{HM}(z) + p_{-z} w_2(-z) - \kappa.$$

Thus, we can rewrite Π_{GP} from (A.21) as follows:

$$\Pi_{GP} := p_A p_B w_{HH} + \sum_{z=A,B} p_{-z}(1 - p_z)\Big\{w_2(-z) + \alpha\Omega(-z) + \lambda_{-z} - \kappa\Big\}.$$

We will use these relations in the following.

A.5.1 Comparative Statics with Respect to δ_z

Consider that $\kappa = 0$. By continuity, the statements hold for sufficiently small $\kappa \geq 0$. Note $p_z w_2(z) = \lambda_z$. This relationship implies

$$w_{HL}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa = w_{HM}(z) + p_{-z} w_2(-z) - \kappa$$

Thus, we can rewrite Π_{GP} from (A.21) as follows:

$$\Pi_{GP} := p_A p_B w_{HH} + \sum_{z=A,B} p_{-z}(1 - p_z)\Big\{w_2(-z) + \alpha\Omega(-z) + \lambda_{-z} - \kappa\Big\}$$

We note that δ_z affects Π_{GP} only via $\alpha\Omega(z)$, where Π_{GP} increases in $\alpha\Omega(z)$. Using the results from Proposition 2, we show that $\Omega(z)$ increases in δ_z.

Note that

$$\frac{d\Omega(z)}{d\delta_z} = (X_z + x - w_2(z) - V_z)\frac{d\theta_z}{d\delta_z} - (1 - \theta_z)\left(\frac{\partial V_z}{\partial \delta_z} + \frac{\partial V_z}{\partial \gamma_z}\frac{\partial \gamma_z}{\partial \delta_z}\right).$$

As $\kappa = 0$, we have $\theta_z = \xi$ and therefore $\frac{d\theta_z}{d\delta_z} = 0$.

When $\gamma_z = 0$, then $\frac{d\Omega(z)}{d\delta_z} = 0$. When $\gamma_z = 1$, then $\frac{\partial \gamma_z}{\partial \delta_z} = 0$, as well as $\frac{\partial V_z}{\partial \delta_z} < 0$. Therefore, $\frac{d\Omega(z)}{d\delta_z} > 0$, which implies $\frac{d\Pi_{GP}}{d\delta_z} > 0$.

Last, consider $\gamma_z \in (0, 1)$. Thus, by Proposition 2, we have:

$$\mathcal{L}_z = (1 - \theta_z)V_z - I + \lambda_z = 0.$$

As $\theta_z = \xi$, this implies $0 = \frac{d\mathcal{L}_z}{d\delta_z} = (1 - \xi)\frac{dV_z}{d\delta_z}$. Thus, $\frac{d\Pi_{GP}}{d\delta_z} = 0$.

It follows that Π_{GP} increases in δ_z and strictly so, whenever $\gamma_z = 1$. A sufficient condition

for $\gamma_z = 1$ is that ξ is sufficiently small.

A.5.2 Comparative Statics with Respect to X_z

We note that X_z affects Π_{GP} only via $\Omega(z)$, where Π_{GP} increases in $\Omega(z)$ when $\alpha = 1$. Suppose $\alpha = 1$. Using the results from Proposition 2, we show that $\Omega(z)$ increases in X_z.

We note that due to $\kappa = 0$, we have $\theta_z = \xi$. Using

$$V_z = \frac{1-\delta_z}{(1-\delta_z)+\delta_z\gamma_z}\left[p_z(X_z - w_2(z))\right],$$

we rewrite

$$\Omega(z) = (1-\xi)(X_z - w_2(z))\left(1 - \frac{(1-\delta_z)p_z}{1-\delta_z+\delta_z\gamma_z}\right) + (1-\xi)x.$$

When $\gamma_z \in \{0,1\}$, then, clearly, $\Omega(z)$ increases with X_z, and so does Π_{GP}.

When $\gamma_z \in (0,1)$, then

$$\mathcal{L}_z = (1-\xi)V_z - I + \lambda_z = 0.$$

Thus, $\frac{d\mathcal{L}_z}{dX_z} = 0$, which implies $(1-\xi)\frac{dV_z}{dX_z} = 0$, which means $\frac{\partial\gamma_z}{\partial X_z} \geq 0$. Since $\frac{(1-\delta_z)p_z}{1-\delta_z+\delta_z\gamma_z}$ decreases in γ_z, it follows that $\Omega(z)$ increases in X_z, and so does Π_{GP}.

A.6 Proof of Corollary 4

Consider that $\kappa = \alpha = 0$. By continuity, the statements hold for sufficiently small $\kappa \geq 0$. Thus, $w_2(z) = \frac{\lambda_z}{p_z}$, i.e., $w_2(A) = \frac{\lambda}{\Delta} + \frac{\sigma}{p}$ and $w_2(B) = \frac{\lambda}{p} - \frac{\sigma}{\Delta}$.

Note $p_z w_2(z) = \lambda_z$. This relationship implies

$$w_{HL}(z) = w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa = w_{HM}(z) + p_{-z}w_2(-z) - \kappa$$

Using this relation as well as the symmetry — where we write $p_z = p$ $\lambda_z = \lambda$, etc., we can rewrite Π_{GP} from (A.21) as follows:

$$\Pi_{GP} := p^2 w_{HH} + p(1-p)\left\{2\left(\frac{\lambda}{p} + \lambda - \kappa\right) + \alpha\Omega_A + \alpha\Omega_B\right\}$$

Due to $\alpha = 0$, we have $\alpha\Omega(z) = 0$.

Further, note that by (A.20), we have:

$$w_{HH} = 2\left[\frac{\lambda}{p} + \lambda + \alpha\Omega(A) + \alpha\Omega(B)\right] - \frac{p}{p}(\lambda - \sigma) - \kappa.$$

Hence, w_{HH} increases in σ. This implies that when α is small, then also Π_{GP} increases in σ.

A.7 Proof of Proposition 7

We consider the postulated crown-jewel contract $\mathcal{C}^{Crown}$ — a second best solution. Here, β indicates whether the contract allows the GP to continue a successful project upon $\alpha = 1$ and double success. As discussed in the main text, this contract stipulates the same payments $w_{HM}(z)$ and $w_{HL}(z)$ as from Proposition 3 (with α set to one). It only modifies the payment for double success, which now becomes w_{HH}^{Crown}. In particular, the crown-jewel contract takes the following form

$$\begin{aligned} w_{HM}(z) &= w_2(z) + \Omega(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}, \\ w_{HL}(z) &= w_2(z) + \Omega(z) + \lambda_z - \kappa, \\ w_{HH}^{Crown} &= \max_{z=A,B} \left\{ w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + (1-\beta)\Omega(z), w_{HH}^{IC} \right\}. \end{aligned} \tag{A.28}$$

Here, $\Omega(z)$ is from (3), determined by γ_z, θ_z, V_z from Proposition 2. We argue below that the continuation stage subgame indeed is characterized through Proposition 2 — i.e., this Proposition also applies in this context.

First, observe that when $\beta = 0$, this contract is identical to the baseline contract from Proposition 3 — with α set to one. Since this case is already treated in Proposition 3, we focus unless otherwise mentioned on $\beta = 1$. With slight abuse of notation, we omit the superscript "$Crown$" when no confusion is likely to arise, i.e., we write $w_{HH}^{Crown} = w_{HH}$.

Consider $\beta = 1$ Note that by construction, the contract $\mathcal{C}^{Crown}$ satisfies incentive constraints (IC-1A), (IC1-B), (IC-2), (IC-3A), (IC-3B), as well as (IC-5), for any $\alpha \in \{0, 1\}$. However, incentive constraint (IC-4) is satisfied only if α happens to equal $\alpha = 0$, while, for $\alpha = 1$, this constraint is violated.

We assume that λ^0 is sufficiently small — akin to assumption 2 — so that $w_{HH} > w_{HH}^{IC}$ and (IC-5) does not bind — in either scenario, i.e., for all α, β.

Crown-Jewel Funds. Suppose $\beta = 1$. Note that $\alpha \in \{0, 1\}$ realizes at time $t = 1$ and is privately observed by the GP. Suppose now both projects have succeeded. If $\alpha = 0$ (with probability $1 - \phi$), the GP derives payoff $w_{HM}(-z) - \kappa + \lambda_z + w_2(z) \leq w_{HH}$ from continuing project z while exiting $-z$ at time $t = 1$. Thus, the GP is better off exiting both projects at time $t = 1$.

If, on the other hand, $\alpha = 1$ (with probability ϕ), the GP derives payoff $w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \Omega(z)$ from continuing project z, while exiting $-z$ at time $t = 1$. This payoff is larger than w_{HH} for at least one $z \in \{A, B\}$. Thus, there exists z such that the GP then continues z and liquidates $-z$, when both projects have succeeded by time $t = 1$.

Note that when continuing a successful project, the GP misrepresents its state to legacy LPs as $s_z = M$.

Adverse Selection at Continuation Stage. We argue that (even under $\beta = 1$) the adverse selection problem at the continuation stage remains analogous to the baseline case, with the solution characterized in Proposition 2. Thus, the same values of θ_z, V_z, γ_z apply.

To make this argument, we conjecture that γ_z and θ_z are equal to their baseline values from Proposition 2. Given this, we argue that new LPs pay the same price V_z as in the baseline. This is equivalent to showing that they hold the same beliefs and possess the same information as in the baseline setting—a point that can be justified by the verbal argument below. Once the price V_z is shown to match the baseline, it follows that γ_z and θ_z must also coincide with their baseline values, thereby confirming the conjecture.

First, note that when the GP continues an H-type project, it misrepresents its type toward legacy LPs as M-type, since the contract does not incentivize truthful revelation in this case. On the other hand, when the GP has an M or L-type project, it truthfully reveals the type to legacy LPs, as stipulated by the contract.

To see this, note that the GP continues a successful project if and only if (i) both projects have succeeded by time 1, and (ii) $\alpha = 1$. When $\alpha = 0$, new LPs are never offered a H-type project, thus they break even under price V_z. When $\alpha = 1$, new LPs may be offered an H-type project, but the GP extracts the full surplus in this case via the side contract. Thus, the new LPs also break even under price V_z. Thus, the price V_z remains unchanged.

Since the price V_z remains unchanged from the baseline, it follows from the expression $\theta_z = \min\left\{\xi + \frac{\kappa}{V_z}, 1\right\}$ that θ_z also remains unchanged. Given equation (A.8), this in turn implies that γ_z is identical to its baseline value.

In sum, we have argued that the continuation stage game—specifically, the quantities θ_z, γ_z, and V_z as characterized in Proposition 2—applies equally in this setting. While a formal treatment could be provided to substantiate this claim, we omit it here for brevity.

Project Choice. Consider now that both projects have succeeded and $\alpha = \beta = 1$. The GP finds it optimal to continue project z rather than $-z$ (while always liquidating the project she does not continue) if:

$$w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \Omega(z) \geq w_{HM}(z) - \kappa + \lambda_{-z} + w_2(-z) + \Omega(-z)$$

With $w_{HM}(z) = w_2(z) + \Omega(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}$, this inequality can be rewritten as:

$$p_{-z}w_2(-z) - \left(p_{-z}w_2(-z) - \lambda_{-z}\right)\mathbb{I}\{\gamma_{-z} > 0\} \geq p_z\, w_2(z) - \left(p_z w_2(z) - \lambda_z\right)\mathbb{I}\{\gamma_z > 0\}. \tag{A.29}$$

We show that $\lambda_z \leq \lambda_{-z}$ implies this inequality — except in the special case that $\kappa > \max\{\lambda_A, \lambda_B\}$ and $\gamma_z > 0 = \gamma_{-z}$. This means that except in this case, $\lambda_{-z} \geq \lambda_z$, implies that continuing project z is optimal. That is, the GP picks the project with the lower agency cost parameter for the crown-jewel fund.

We distinguish several cases:

1. First, suppose that $\kappa < \min\{\lambda_A, \lambda_B\}$, so that $p_z w_2(z) = \lambda_z$. In this case, the inequality (A.29) reduces to $\lambda_{-z} \geq \lambda_z$ (which holds by assumption).

2. Second, suppose that $\kappa > \max\{\lambda_A, \lambda_B\}$, so that $p_z w_2(z) = \kappa$. The inequality (A.29) then becomes
$$-\big(\kappa - \lambda_z\big)\,\mathbb{I}\{\gamma_z > 0\} \geq -\big(\kappa - \lambda_{-z}\big)\,\mathbb{I}\{\gamma_{-z} > 0\}.$$
If both γ_z, γ_{-z} are strictly bigger than zero, the inequality becomes $\lambda_{-z} \geq \lambda_z$, as desired. If γ_z, γ_{-z} are both zero or $\gamma_{-z} > 0 = \gamma_z$, the inequality holds too.

3. Third, suppose $\lambda_z \leq \kappa \leq \lambda_{-z}$. Thus, $p_z w_2(z) = \kappa \leq p_{-z} w_2(-z) = \lambda_{-z}$. The inequality (A.29) becomes:
$$\lambda_{-z} - \kappa \geq -\big(\kappa - \lambda_z\big)\,\mathbb{I}\{\gamma_z > 0\},$$
which holds under the stated parameter conditions, since the right-hand-side is weakly negative and the left-hand-side is weakly positive.

Payoff Changes. We calculate the payoff changes when β changes from 0 to 1, where Δ denotes the induced change with respect to β. Suppose z is picked for crown-jewel fund. Relative to $\beta = 0$, setting $\beta = 1$ changes NPV by

$$\Delta NPV = -p_A p_B \phi (I - x - \lambda_z).$$

Indeed, in case of double success — which happens with probability $p_A p_B$, the GP inefficiently continues the successful project with probability ϕ, when $\alpha = 1$. Upon continuation of an H-type project, additional cash flow x and private benefits λ_z are generate, but the cost is I. Generally, the deadweight loss from inefficient continuation is $I - x - \lambda_z$.

Raising β from 0 to 1 changes GP payoff by

$$\Delta \Pi_{GP} = -p_A p_B (1 - \phi) \Omega(z).$$

This is because the GP's payment for double success — which happens with probability $p_A p_B$ — is now lowered (relative to $\beta = 0$) by $\Omega(z)$. The GP collects this payment when $\alpha = 0$ (with probability $1 - \phi$). Otherwise, when $\alpha = 1$, the GP can extract an additional payoff of $\Omega(z)$, i.e., in this contingency, its payoff does not change relative to the $\beta = 0$

scenario. Importantly, recall that new LPs will realize zero payoff irrespective of α — they merely break even.

Thus, raising β from 0 to 1 changes legacy LPs' payoff by:

$$\Delta\Pi_{LP} = \Delta NPV - \Delta\Pi_{GP} - \Delta\Pi_{New} = p_A p_B \left[(1-\phi)\Omega(z) - \phi(I - x - \lambda_z)\right].$$

Using $\Omega(z) = (1-\theta_z)(X_z + x - w_2(z) - V_z)$ and

$$V_z = \frac{(1-\delta_z)p_z}{1-\delta_z+\delta_z\gamma_z}(X_z - w_2(z))$$

from (A.2), we obtain that legacy LPs are better off with $\beta = 1$ (i.e., $\Delta\Pi_{LP} \geq 0$) when

$$p_A p_B(1-\phi)\Omega(z) \geq p_A p_B \phi \left[(I - x - \lambda_z)\right].$$

We can rewrite this inequality as follows:

$$(1-\theta_z)\left\{(X_z - w_2(z))\Big[1 - \frac{(1-\delta_z)p_z}{(1-\delta_z)+\delta_z\,\gamma_z}\Big] + x\right\} \geq \frac{\phi(I - x - \lambda_z)}{1-\phi}, \tag{A.30}$$

where $\theta_z = \min\left\{\xi + \frac{\kappa}{V_z}, 1\right\}$. This expression coincides with (8), as desired. We recall that $\theta_z = \xi + \frac{\kappa}{V_z}$.

A.8 Alternative Structure for Crown-Jewel Funds

We consider the alternative crown-jewel contract $\mathcal{C}^{Crown}$ — a second best solution. Here, $\beta \in \{0,1\}$ as to whether the contract the GP to continue a successful project (while misrepresenting as M-type) when $\alpha = 1$. The altnernative crown-jewel contract takes the following form

$$\begin{aligned} w_{HM}(z) &= w_2(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\} + (1-\beta)\Omega(z), \\ w_{HL}(z) &= w_2(z) + \lambda_z - \kappa + (1-\beta)\Omega(z), \\ w_{HH} &= \max_{z=A,B}\left\{w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + (1-\beta)\Omega(z), w_{HH}^{IC}\right\}. \end{aligned} \tag{A.31}$$

Here, $\Omega(z)$ is from (3), determined by γ_z, θ_z, V_z from Proposition 2. We argue below that the continuation stage subgame indeed is characterized through Proposition 2 — i.e., this Proposition also applies in this context.

First, observe that when $\beta = 0$, this contract is identical to the baseline contract from Proposition 3 — with α set to one. Since this case is already treated in Proposition 3, we focus unless otherwise mentioned on $\beta = 1$.

Consider $\beta = 1$ Note that by construction, the contract $\mathcal{C}^{Crown}$ satisfies incentive constraints (IC-1A), (IC1-B), (IC-2), (IC-3A), (IC-3B), (IC-4) and (IC-5) — that is, all incentive constraints — if and only if $\alpha = 0$ (with probability $1 - \phi$).

We assume that (IC-5) does not bind in line with previous analysis, i.e., λ^0 is not too large. That is, $w_{HH} > w_{HH}^{IC}$ throughout.

Crown-Jewel Funds. Suppose $\beta = 1$. Under the proposed contract, the GP follows the prescribed strategy when $\alpha = 0$ (with probability $1-\phi$) and exits successful projects at time $t = 1$. However, with probability ϕ, we have $\alpha = 1$ and the GP continues the successful project, whenever given the opportunity. Note that when the GP continues an H-type project, it misrepresents its type toward legacy LPs as M-type, since the contract does not incentivize truthful revelation in this case. On the other hand, when the GP has an M or L-type project, it truthfully reveals the type to legacy LPs, as stipulated by the contract.

Adverse Selection at Continuation Stage. We argue that the adverse selection problem at the continuation stage remains analogous to the baseline case, with the solution characterized in Proposition 2. Thus, the same values of θ_z, V_z, γ_z apply. When $\alpha = 1$, the GP can side-contract to capture the incremental value in the H-type state, leaving new LPs with expected payoff exactly V_z. Hence their break-even price remains V_z even though H may be (rarely) continued.

As before (see Proof of Proposition 7 for details), this follows from the fact that irrespective of α, new LPs break even under the price V_z. Given the price V_z, θ_z and γ_z are determined as in the baseline.

This determines $\Omega(z)$.

Project Choice. First, if only one project has been successful and $\alpha = 1$, the GP continues the successful one, so there is no choice.

Second, consider that both projects have succeeded and $\alpha = 1$. The GP finds it optimal to continue project z rather than $-z$ (while always liquidating the project she does not continue) if:

$$w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \Omega(z) \geq w_{HM}(z) - \kappa + \lambda_{-z} + w_2(-z) + \Omega(-z)$$

With $w_{HM}(z) = w_2(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\}\big\}$, this inequality simplifies to:

$$\begin{aligned} p_{-z}w_2(-z) \; - \; & \big(p_{-z}w_2(-z) - \lambda_{-z}\big)\,\mathbb{I}\{\gamma_{-z} > 0\} + \Omega(z) \\ \geq p_z\, w_2(z) - \; & \big(p_z w_2(z) - \lambda_z\big)\,\mathbb{I}\{\gamma_z > 0\} + \Omega(-z). \end{aligned}$$

Payoff Changes. We now assess how payoffs and NPV change upon raising β from 0 to 1. We denote this induced change by Δ. Suppose (without loss of generality) that A is picked for crown-jewel fund in case the GP has a choice and both projects have succeeded by time $t = 1$.

Raising β from 0 to 1 changes NPV by:

$$\begin{aligned}\Delta NPV := -\phi\Big\{ & p_A p_B (I - x - \lambda_A) \\ & + p_A(1-p_B)\Big[I - x - \lambda_A - \delta_B\gamma_B(I-\lambda_B) + (1-\delta_B)(p_B X_B - I)\Big] \\ & + p_B(1-p_A)\Big[I - x - \lambda_B - \delta_A\gamma_A(I-\lambda_A) + (1-\delta_A)(p_A X_A - I)\Big]\Big\}.\end{aligned}$$

To understand that change, note that the inefficient continuation of an H-type project only arises when $\alpha = 1$, which happens with probability ϕ. In case of double success (with probability $p_A p_B$), the inefficient continuation generates additional cash flow x and private benefits for the GP λ_A, but has a cost I. Generally, the deadweight loss from inefficient continuation is $I - x - \lambda_z$.

In case of single success (with probability $p_z(1-p_{-z})$), the GP continues the successful project z, reducing NPV by $I - x - \lambda_z$. If, in addition, project $-z$ remains high-potential (with conditional probability $1 - \delta_{-z}$), the NPV from continuing high-potential projects, i.e., $p_{-z}X_{-z} - I$ is lost too. By contrast, if $-z$ has failed (with conditional probability δ_{-z}), the failed project is at least not continued, saving deadweight loss from continuing a failed project, $(I - \lambda_{-z})$.

Thus, raising β from 0 to 1 changes the GP's payoff by

$$\Delta\Pi_{GP} = -(1-\phi)\Big[p_A(1-p_B)\Omega(A) + p_B(1-p_A)\Omega(B) + p_A p_B(\Omega(A)+\Omega(B))\Big] - \phi p_A p_B \Omega(B).$$

To understand this payoff change, note that, with probability $1 - \phi$, we have $\alpha = 0$ and the GP cannot extract surplus via a side deal. This lowers the payment when z succeeds and $-z$ does not by $\Omega(z)$ — this event occurs with probability $p_z(1-p_{-z})$. When $\alpha = 1$ and one project succeeds, the GP's payoff is independent of β, because when the other project is failed, the contract makes the GP indifferent between continuing and stopping, so the change in β has no effect on its payoff in this contingency.

When both projects succeed, then we can derive — analogous to (A.20) — that:

$$\begin{aligned}w_{HH} = & \big[w_2(A) + w_2(B) - \kappa + \lambda_A + \lambda_B + (1-\beta)\big[\Omega(A)+\Omega(B)\big]\big] \\ & - \min_{z=A,B}\Big[p_z w_2(z) - (p_z w_2(z) - \lambda_z)\mathbb{I}\{\gamma_z > 0\}\Big].\end{aligned}$$

Thus, raising β from 0 to 1 lowers w_{HH} by $\Omega(A) + \Omega(B)$. With probability $1 - \phi$, the GP exits and collects this payment. With probability ϕ, the GP exits B, continues A, engages in a side deal, and overall collects

$$w_{HM}(B) + \lambda_A - \kappa + w_2(A) + \Omega(A) = w_{HH} + \Omega(A)$$

where w_{HH} is under $\beta = 1$.

Now calculate legacy LPs' payoff change as

$$\Delta\Pi_{LP} = \Delta NPV - \Delta\Pi_{GP}.$$

LPs are better off with $\beta = 1$ when $\Delta\Pi_{LP} \geq 0$. This inequality always holds when ϕ is sufficiently small, as $\Delta NPV = 0$ for $\phi = 0$ while $\Delta\Pi_{GP} < 0$ for $\phi = 0$. The same trade-offs as those highlighted in Proposition 7, but the expressions are significantly more complex.

A.9 Proof of Corollary 5

In all scenarios, suppose that $-z$ has succeeded and the GP considers continuing z.

Case 1. $\xi \to 1$. First, suppose that $\xi \to 1$, which implies $\theta_z \to 1$. Thus, (8) does not hold (for any $\phi > 0$) and legacy LPs write the baseline contract from Proposition 3. Moreover, we have $\gamma_z = 0$ according to (A.8). Hence, if project z is continued, it of type M with probability one. This readily implies an expected cash flow of $p_z X_z$.

Case 2. $\xi \to 0$. Second, suppose that $\xi \to 0$. Then (A.8) implies $\gamma_z \geq 0$, where the inequality is strict and γ_z if κ is not too alrge. By assumption, $I + x - \lambda_z$ is not too large, i.e., $I - \lambda_z \to 0$, so the limit $\xi \to 0$ implies $\theta_z < 1$ and therefore (8) holds. Thus, if both project have been successful, then project z is continued.

With probability p_z, project z has been successful in the legacy fund. It is then continued with probability ϕ. With probability $(1-p_z)\delta_z\gamma_z$, project z has failed and is continued. With probability $(1-p_z)(1-\delta_z)$ project z is continued as M-type. With probability $(1-p_z)\delta_z(1-\gamma_z)$, project z has failed and is not continued.

Thus, the only scenarios in which z is not continued is (i) success and probability $1-\phi$ and (ii) failure —- with probability $(1 - p_z)\delta_z$ — and the randomization yields no continuation, which happens with probability $(1 - \gamma_z)$. Therefore, we denote by

$$q_z := 1 - (1 - p_z)\delta_z(1 - \gamma_z) - p_z(1 - \phi)$$

the probability of continuing project z.

Conditional on continuation, project z is of H-type with probability $q_z^H := \frac{p_z\phi}{q_z}$. It is of M-type with probability $q_z^M := \frac{(1-p_z)(1-\delta_z)}{q_z}$, and it is of L-type with probability $\frac{(1-p_z)\delta_z\gamma_z}{q_z}$. Expected cash flow equals (conditional on continuation):

$$\mathcal{E} := \left[q_z^M p_z + q_z^H\right] X_z + q_z^H x.$$

We can calculate

$$\mathcal{E} - p_z X_z = \frac{p_z}{q_z}\left(\phi - \delta_z\gamma_z\right)\left(1 - p_z\right) X_z + q_z^H x,$$

which is positive for $\phi > \delta_z\gamma_z$.

B Other Results

B.1 Long-Term Contracting and Proof of Proposition 9

We sketch how solve the long-term contracting problem. Legacy LPs sign a long-term contract $\mathcal{C}^{LT}$ with the GP. That is, $\xi = 1$. The long-term contract takes the form

$$\mathcal{C}^{LT} = (w_{HM}(z), w_{HL}(z), w_{HH}, w_2(z)).$$

As there is no asset sale, we have $\Omega(z) = \gamma_z = 0$ under the long-term contract.

The long-term contract is incentive compatible as long as (IC-1A), (IC1-B), (IC-2), (IC-3A), (IC-3B), (IC-4),(IC-5), and (IC-C) are satisfied. These incentive constraint ensure truth telling of the state at $t = 1$ and full effort. They also ensure that successful and failed projects are exited at $t = 1$.

The GP's payoff equals

$$\Pi_{GP} := p_A p_B w_{HH} + \sum_{z=A,B} p_{-z}(1-p_z)\Big\{(1-\delta_z)[p_z w_2(z) - \kappa + w_{HM}(-z)] + \delta_z w_{HL}(-z)\Big\}.$$

Thus, the GP collects w_{HH} in case of double success. It collects $w_{HM}(z)$ is z is exited successfully at $t = 1$, and $-z$ is continued. It collects $w_{HL}(z)$ if z is exited successfully at time $t = 1$, while $-z$ is also exited but not successfully.

The NPV under the optimal contract is given by:

$$NPV = p_A p_B (X_A + X_B) + \sum_{z=A,B} p_{-z}(1-p_z)\Big\{X_{-z} + (1-\delta_z)(p_z X_z - I)\Big\}.$$

The long-term contract is designed to maximize LP payoff:

$$\Pi_{LP} = NPV - \Pi_{GP},$$

subject to all relevant incentive and participation constraints. Thus, it is optimal to minimize agency rents, as the contract does not directly affect NPV.

It is clear that setting the minimum payment $w_2(z)$ subject to (PC-C) and (IC-C) is optimal to minimize agency rent, leading to the continuation fund contract from Proposition 1. We note that only M-type projects are continued.

From then, one easily verifies that the payments from Proposition 3, since these are the minimum payments that satisfy all incentive constraints, given $w_2(z)$. Since $w_2(z)$ remains identical to the baseline, the other payments do too. Formally, this can be shown repeating the arguments in the proof of Proposition 3. Thus, the optimal long-term contract payments are $w_2(z)$ from Proposition 1, with the remaining payments obtained from Proposition 3 upon setting $\alpha\Omega(z) = \gamma_z = 0$, as well as $\xi = \theta_z = 1$.

B.2 Continuation Vehicles of Continuation Vehicles

We extend the model to allow for continuation vehicles of continuation vehicles, extending the model for an additional period from $t = 2$ to $t = 3$. The model until time $t = 2$ remains identical. The following assumptions are also intended to stay within our setting as much as possible and to design the continuation from $t = 2$ to $t = 3$ analogously to the one from $t = 1$ to $t = 2$ — featuring the identical moral hazard problem. However, there is no more coinvestment by the GP in the second CV, for simplicity.

Suppose project z is continued from $t = 1$ to $t = 2$, where firm z has firm type $s_z \in \{L, M, H\}$ at time $t = 1$. Then, at time $t = 2$, the GP privately observes both the price at which it can be sold to external buyers and the state of the firm — which we denote by $s_{z,2} \in \{L, M, H\}$. That is, s_z denotes firm type at $t = 1$, while $s_{z,2}$ denotes firm type at $t = 2$. The GP can then decide whether to continue firm z for another period from $t = 2$ to $t = 3$ — which, however, requires another investment of I dollars. Regarding continuation and firm states, we make the following assumptions.

First, if an L-type firm was continued from time $t = 1$ to $t = 2$ — that is, $s_z = L$ — the firm remains in state L and has permanently a price or value of zero to external buyers. Hence, $s_z = L$ implies $s_{z,2} = L$.

Second, if M-type firm was continued from time $t = 1$ to $t = 2$ — that is, $s_z = M$ — the firm either succeeds at time $t = 2$, which happens with probability p_z (assuming the GP works). In case success is achieved by time $t = 2$ yet has not been achieved at time $t = 1$, the firm can be sold at price X_z at time $t = 2$ or at price $X_z + x$ at time $t = 3$. Otherwise, if

no success is achieved (with probability $1 - p_z$), the firm remains in state M, that is, there is no more failure risk at time $t = 2$, for simplicity. The firm then succeeds by time $t = 3$ with probability p_z, yielding resale value of X_z; the resale value is zero otherwise.

Third, if H-type firm is continued, this firm can be sold at either time $t = 2$ or time $t = 3$ at price $X_z + x$. Put differently, as before, an H-type firm remains in that state. Crucially, the price no more improves after $X_z + x$ is reached. This value is reached precisely one period after the success — that is, if a project has succeeded by time t, its resale value at time t is X_z, while it is $X_z + x$ at times $t + 1$ and later.

The state $s_{z,2}$ at time $t = 2$ is privately observed by the GP. Likewise, s_z is privately observed by the GP at time $t = 1$.

Second Continuation Fund Contract. The second continuation fund stipulates payments to the GP of $w_3(z)$ in case of success (a cash flow of X_z or higher) and zero otherwise. Analogously to before, this payment is such that the GP has incentives to work — in case project z is of type M — and participates at fair terms. This leads to

$$w_3(z) = \frac{\lambda_z}{p_z}.$$

We note that the provided the project is of type M at time $t = 2$, the GP is always better off raising the second continuation fund, as opposed to stopping. If a payment of $X_z + x$ realizes, the GP is also paid $w_3(z)$.

Note there is no more coinvestment by the GP in the second continuation fund, by assumption.

First Continuation Fund Contract. The contract in the first continuation fund pays the GP $w_2(z)$ in case project z is successfully exited at time $t = 2$, producing a cash flow of at least X_z. It pays $w^2_{HM}(z)$ in case z is continued until $t = 3$. However, when z is exited at price zero at $t = 2$, then the payment is zero — the reason is that the fund has no cash flows out of which payments can be made. That is, we assume payments must be made from fund cash flows — zero cash flow implies zero payment.

We already note that there is no point in paying the GP for continuing — which just exacerbates incentive problems (specifically, it creates incentives to continue bad, failed projects). Thus, we can already set $w^2_{HM}(z) = 0$.

We assume that all legacy LPs that were present at $t = 0$ exit at time $t = 2$. These legacy LPs had an information advantage vis-a-vis new LPs. Thus, in the second continuation vehicle, we have both new LPs entering at time $t = 2$ and the LPs that entered at $t = 1$. As we argue, these two types of LPs have the same information sets.

Note that when the GP continued a failed project in the first continuation fund and this project does not succeed by $t = 2$, the GP cannot be paid for terminating at price zero. As discussed before, GP is paid out of cash flows and its payment must be smaller than the total cash flow. continues projects after failure, so generally $\eta_{z,2}(s_{z,2}) > 0$ for $s_{z.2} \in \{L, M\}$. Under these circumstances, it is not possible to incentivize the GP to disclose the failure and to stop failed projects at $t = 2$ (unless the market breaks down and below condition (B.33) does not hold). Therefore, provided it is feasible, the GP always continues a failed project until $t = 3$. That is, if the GP continues a project with type $s_z = L$ at time $t = 1$, then it holds on to the project at $t = 3$. This implies that potentially new LPs entering at time $t = 2$ and those that entered at time $t = 1$ have the same information set around the firm's type $s_{z,2}$.

As in the baseline, we focus on incentive-compatible legacy fund (primary fund) contracts $\mathcal{C}^{Leg}$ that induce the GP to exit successful projects, full effort, and truthful reporting vis-a-vis legacy LPs.

Valuation and Beliefs at $t = 2$. New LPs entering at time $t = 2$ and those that entered at time $t = 1$ have the same information set around the firm's type $s_{z,2}$ — which in equilibrium is either L or M (as no H-type projects are continued).

The fair value (gross of investment) of the second continuation fund under investor beliefs $\eta_{z,2}(s_{z,2})$ equals:

$$V_{z,2} = \eta_{z,2}(M)p_z(X_z - w_3(z)). \tag{B.32}$$

We note that because we focus on contracts that terminate successful projects as soon as possible, we have $\eta_{z,2}(H) = 0$, i.e., in equilibrium, no high type projects are carried over.

Raising a second continuation fund is feasible as long as

$$V_{z,2} - I > 0 \tag{B.33}$$

Otherwise, LPs are unwilling to provide the capital outlay required. Recall that there is no more coinvestment by the GP in the second continuation fund (by assumption).

The probability with which project z was L-type at time $t = 1$ equals $\frac{\delta_z\gamma_z}{1-\delta_z+\delta_z\gamma_z}$ where γ_z is endogenous. Thus, given these beliefs, the probability of success by time $t = 2$ equals

$$\frac{p_z(1-\delta_z)}{1-\delta_z+\delta_z\gamma_z}.$$

When new LPs see that z did not succeed by time $t = 2$ and z is continued in a second

continuation fund until $t = 3$, Bayes' rule implies:

$$\eta_{z,2}(L) = \frac{\delta_z \, \gamma_z}{(1 - p_z)\,(1 - \delta_z) + \delta_z \, \gamma_z}$$

and

$$\eta_{z,2}(M) = 1 - \eta_{z,2}(L) = \frac{(1 - p_z)\,(1 - \delta_z)}{(1 - p_z)\,(1 - \delta_z) + \delta_z \, \gamma_z}.$$

Continuation Stage Game at $t = 2$. Take θ_z — the share of equity in the first CV held by legacy LPs as given.

We assume legacy investors (who were there at $t = 0$) exit and are replaced by new investors, arriving at $t = 2$. Further, fraction $(1 - \xi_2)$ of time-1 LPs exit. Both sell at price $V_{z,2} - I$, analogous to our previous modeling Thus, there is an overall payout $\ell_{z,2}$ that must be raised from new investors, which satisfies $\ell_{z,2} = \big[\theta_z + (1 - \theta_z)(1 - \xi_2)\big](V_{z,2} - I)$. Further, new LPs, who arrived at $t = 1$, contribute $(1 - \theta_z)\xi_2 I$ units of new capital,.

This means the GP must raise

$$I + \big[\theta_z + (1 - \theta_z)(1 - \xi_2)\big](V_{z,2} - I) - (1 - \theta_z)\xi_2 I.$$

from new investors/LPs who arrive at $t = 2$. Above equation can be simplified to $\big[1 - (1 - \theta_z)\xi_2\big]V_{z,2}$. New time-2 investors break even given their beliefs and acquire stake $1 - \theta_{z,2}$ in the second continuation fund at price $V_{z,2}$. Thus:

$$(1 - \theta_{z,2})V_{z,2} = \big[1 - (1 - \theta_z)\xi_2\big]V_{z,2}.$$

As such, we can solve for the stake in the second CV that is retained by time-1 LPs, $\theta_{z,2}$. This stake satisfies

$$\theta_{z,2} = (1 - \theta_z)\,\xi_2, \tag{B.34}$$

where $1 - \theta_z$ was the stake that time-1 LPs held in the first (!) CV.

As in the baseline, when the project is of high quality, the GP can side-contract with new LPs, which allows the GP to extract surplus $\alpha\Omega_2(z)$ where:

$$\Omega_2(z) = (1 - \theta_{z,2})\,[X_z + x - w_3(z) - V_{z,2}]\,,$$

and $\alpha \in \{0, 1\}$. The idea is as in the baseline. When the GP deviates and continues a H-type project, then exiting LPs sell too cheaply. They sell $1 - \theta_{z,2}$ shares (out of a unit mass) at price $V_{z,2}$ to new LPs. However, in reality, a share is worth $X_z + x - w_3(z)$. The GP can extract this surplus from new LPs via a side deal.

Incentive Constraints in First Continuation Fund. The payment $w_2(z)$ must satisfy — by assumption — (PC-C). The remaining incentive constraints change.

First, consider z has succeeded by time $t = 2$. For the GP to have incentives to exit z at time $t = 2$ as opposed to at time $t = 3$, we have:

$$w_2(z) \geq w_3(z) + \lambda_z + \alpha\Omega_2(z). \tag{B.35}$$

Second, an incentive constraint regarding effort applies:

$$p_z w_2(z) + (1 - p_z)p_z w_3(z) \geq \lambda_z + p_z w_3(z).$$

which becomes

$$w_2(z) - p_z w_3(z) \geq \frac{\lambda_z}{p_z}. \tag{B.36}$$

Note that when the GP shirks at $t = 1$ for an M-type project, the project remains in state M and will be continued in the second continuation fund where the GP under effort earns $p_z w_3(z)$ in expectation.

We focus on the minimum payment that ensures incentive compatibility and participation constraints, as in the baseline. This yields:

$$w_2(z) = \max\left\{\frac{\kappa}{p_z}, \frac{\lambda_z}{p_z} + p_z w_3(z), w_3(z) + \lambda_z + \alpha\Omega_2(z)\right\}. \tag{B.37}$$

In general, we have $w_2(z) > \max\left\{\frac{\kappa}{p_z}, \frac{\lambda_z}{p_z}\right\}$, i.e., the payment is larger than it was in the baseline.

First Continuation Stage Game. The continuation stage game is analogous to the baseline, but will be affected by the possibility of a second continuation fund. The unique Perfect Bayesian Nash Equilibrium under the intuitive criterion — whose existence and uniqueness we do not formally prove — is described by the following equations.

First, as in the baseline, for the first continuation fund, fraction $1 - \theta_z$ of shares are held by new LPs while legacy LPs retain

$$\theta_z = \min\left\{\xi + \frac{\kappa}{V_z}, 1\right\}. \tag{B.38}$$

Second, the fair price that new investors pay satisfies:

$$V_z = \frac{1 - \delta_z}{(1 - \delta_z) + \delta_z\gamma_z}\left[p_z(X_z - w_2(z))\right] + \left(1 - \frac{(1 - \delta_z)p_z}{(1 - \delta_z) + \delta_z\gamma_z}\right)\left[V_{z,2} - I\right]^+ \tag{B.39}$$

where $[\cdot]^+ = \max\{0, \cdot\}$ accounts for the possibility that raising a second continuation fund may not be feasible.

The second term of (B.39) is new relative to the baseline. Here,

$$\left(1 - \frac{(1-\delta_z)p_z}{(1-\delta_z) + \delta_z \gamma_z}\right)$$

represents the probability with which the continued project z does not succeed in the first continuation fund under equilibrium beliefs. Then, a second continuation vehicle is raised and new LPs who arrive at $t = 1$ realize a value of $V_{z,2} - I$ (provided this is positive) per unit of share they own — either they sell at this price or hold an equity claim that they value at this price.

Next, when a project has failed by time $t = 1$ (i.e., $s_z = L$), the continuation decision of the coalition is as in the baseline, following (A.8) which determines γ_z. However, the payoff from continuation a failed project now changes to

$$\mathcal{L}_z = (1-\theta_z)V_z - I - \kappa + \lambda_z + \mathbb{I}\{V_{z,2} - I \geq 0\}\left[\theta_z(V_{z,2} - I) + \lambda_z\right] \tag{B.40}$$

The term $\left[\theta_z(V_{z,2} - I) + \lambda_z\right]$ stems from the fact that the failed project is rolled on until $t = 3$, giving legacy LPs a payoff of $\theta_z(V_{z,2} - I)$ from selling their share in the second continuation fund and yielding additional private benefits λ_z for the GP. Note that raising the second continuation fund is only feasible when $V_{z,2} - I$, which we account for via the indicator $\mathbb{I}\{V_{z,2} - I \geq 0\}$. Note that unlike in the second CV, there is GP coinvestment in the first CV.

Legacy Fund Contract The legacy fund contract continues to take the form from Proposition 3. We do not re-derive its structure. As shown in Proposition 3 payments satisfy:

$$\begin{aligned}
w_{HM}(z) &= w_2(z) + \alpha\Omega(z) + \lambda_z - p_{-z}w_2(-z) + (p_{-z}w_2(-z) - \lambda_{-z})\mathbb{I}\{\gamma_{-z} > 0\},\\
w_{HL}(z) &= w_2(z) + \alpha\Omega(z) + \lambda_z - \kappa,\\
w_{HH} &= \max_{z=A,B}\left\{w_{HM}(-z) - \kappa + \lambda_z + w_2(z) + \alpha\Omega(z), w_{HH}^{IC}\right\},
\end{aligned}$$

The possibility to raise a second continuation fund affects the legacy contract via γ_A, γ_B and the payment $w_2(A), w_2(B)$. The payment $w_2(z)$ is now higher than in the baseline, overall leading to higher payouts too.

Take-away and Results. The possibility of the second continuation fund has several consequences.

First, we note that $w_3(z)$ increases $w_2(z)$, thereby affecting the continuation fund contract and leading to higher payments relative to the baseline. This tightens incentive constraints in the legacy fund, effectively exacerbating agency issues. As a result, the legacy contract then also features higher payments.

Second, the continuation stage game is affected as follows. First, as shown in (B.39), the valuation for the first continuation fund becomes higher in case a second continuation fund is raised. This reduces θ_z and increases $1-\theta_z$, which increases $\Omega(z)$ and exacerbates agency issues.

Third, $\mathcal{L}_z$ from (B.40) becomes larger than in the baseline, which then introduces incentives to continue failed projects and to raise zombie continuation funds.

Taken together, the possibility to raise a second continuation fund exacerbates agency issues and can increase the propensity to raise zombie funds, as it allows longer roll-on of failed projects.

Proof of Proposition 8 Under symmetry we drop the index z unless confusion is likely to arise, and let p denote the success probability and $\lambda > 0$ the private benefit from shirking. With a second continuation fund allowed, $\alpha = 0$ and κ small, the optimal payment in the second CV is $w_3 = \frac{\lambda}{p}$. The first-CV payment w_2 is given by the maximum of the participation constraint and the binding incentive constraints (see Appendix B.2):

$$w_2 = \max\left\{\frac{\kappa}{p},\ \frac{\lambda}{p} + p\, w_3,\ w_3 + \lambda\right\} = \max\left\{\frac{\kappa}{p},\ \frac{\lambda}{p} + \lambda\right\}.$$

where we substituted $w_3 = \lambda/p$. In the baseline (no second CV), we have from Proposition 1 that $w_2^{\text{base}} = \max\left\{\frac{\kappa}{p},\ \frac{\lambda}{p}\right\}$. Thus, when κ is small, then $w_2 > w_2^{\text{base}}$. This proves the first claim.

For the second claim, we check each payment in the legacy fund:

- For a single success with the other project exited at zero, the GP receives in the legacy fund $w_{HL} = w_2 + \lambda - \kappa > 0$ (recall κ is small). With w_2 strictly higher than in the baseline, this payment w_{HL} is strictly larger than the corresponding baseline payment.

- For a success while the other project is continued, the legacy fund payment equals

$$w_{HM} = w_2 + \lambda - p\, w_2 + \big(p\, w_2 - \lambda\big)\, \mathbb{I}\{\gamma > 0\}.$$

  If $\gamma > 0$, then $w_{HM} = w_2$; if $\gamma = 0$, then $w_{HM} = (1-p)w_2 + \lambda$. In either case w_{HM} is strictly increasing in w_2, hence strictly higher.

- For both projects successful, the payment is

$$w_{HH} = w_{HM} - \kappa + \lambda + w_2.$$

where we used that λ^0 is small in that $w_{HH} > w_{HH}^{IC}$. As both w_{HM} and w_2 are strictly larger than in the baseline, it follows that w_{HH} is strictly larger too than in the baseline.

B.3 Modified Continuation Fund Participation Constraint

The baseline analysis has assumed that condition (PC-C) holds—namely, that the GP acquires its stake in the continuation fund on fair terms under the assumption that $s_z = M$. When (PC-C) binds, this means the GP buys into the continuation at a price consistent with the firm being of type M.

In equilibrium, however, continuation may occur for both L- and M-type firms. As a result, the transfer price may differ from the one implied by $s_z = M$ (with probability one). To capture this explicitly, we could assume instead that the GP acquires its stake in the continuation fund on terms that are fair from the perspective of the new LPs. This would lead to the modified participation constraint:

$$\left(\frac{(1-\delta_z)p_z}{1-\delta_z+\delta_z\gamma_z}\right) w_2(z) \geq \kappa, \tag{B.41}$$

where the term in parentheses reflects the new LPs' expected success probability in the continuation stage. All other elements of the setup remain unchanged.

The continuation fund payment $w_2(z)$ is then determined as the minimum amount satisfying both (IC-C) and (B.41). When (B.41) binds, an important effect emerges: the GP's continuation fund stake $w_2(z)$ increases with $\delta_z\gamma_z$, where δ_z is the probability of an L-type and γ_z is the pooling probability.

This has two key interpretations. First, a lower transfer price allows the GP to coinvest on more favorable terms, strengthening incentives to transfer high-quality (H-type) projects into the continuation fund. Second, because the GP's carried interest resembles a call option— with the hurdle rate analogous to the strike price—a lower transfer price reduces the hurdle relative to the entry valuation, thereby increasing the value of the GP's option-like compensation. This effect is especially strong when the option is already *in the money*, as for H-type projects.

A similar mechanism underlies Corollary 3, which showed that the GP may benefit from higher risk. One could re-prove that result under the modified participation constraint (B.41) and obtain the same conclusion.

B.4 Side Contracts — Robustness

We now argue that one could equally allow the M-type project to produce cash flow $X_z + x$ upon successful continuation without affecting the GP's ability to extract surplus $\Omega(z)$ when $\alpha = 1$.

Consider an M-type project. Suppose that, conditional on success (with probability p_z), the M-type project produces cash flow X_z with probability $1-\chi_z$ and cash flow X_z+x with probability χ_z. Thus, the baseline obtains for $\chi_z = 0$.

Then, a time-2 cash flow of $X_z + x$ is consistent with both M- and H-type projects, i.e., it would not reveal a deviation from the GP. How can the GP now extract surplus via a side contract?

Suppose the GP would like to continue a H-type project and would like to extract surplus $\Omega(z)$ from new LPs via a private side contract. Consider $\alpha = 1$. The GP could propose a side contract that stipulates a time-2 payment of $\tilde{w}_2(z)$ in case of cash flow $X_z + x$ and a payment of $\hat{w}_2(z) \in (0, w_2(z)]$ in case of X_z; 0 otherwise. Thus, if the project returns cash flow X_z only, the GP pays the new LPs $-\hat{w}_2(z)$, while otherwise the new LPs pay the GP $\tilde{w}_2(z)$. The payment from the GP to new LPs is lower than $w_2(z)$ — the reward from the contract $\mathcal{C}_z^{CV}$ in case the continued project produces cash flow of at least X_z. Note that in case the continued project is of type H, then X_z+x with certainty at $t = 2$, and the contract pays $\tilde{w}_2(z)$.

As long as

$$\chi_z \tilde{w}_2(z) + (1-\chi_z)\hat{w}_2(z) < 0,$$

the GP would only offer such contract if it indeed has an H-type project at hand. If the GP had an M-type project, the GP would loose by offering the contract. If the GP had an L-type project, the contract is irrelevant, i.e., never pays. That is, by offering such contract $\chi_z \tilde{w}_2(z) + (1-\chi_z)\hat{w}_2(z) < 0$ to new LPs, the GP essentially could signal the project H-type credibly.

Let us solve for a side contract that allows the GP to extract rents $\Omega(z)$ from new LPs when continuing a H-type project. This can, for instance, be achieved upon setting $\tilde{w}_2(z) = \Omega(z)$ and $\hat{w}_2(z) = -w_2(z)$. This contract is feasible — i.e., only the H-type would offer it — as long as $\chi_z \Omega(z) < (1-\chi_z) w_2(z)$, i.e., as long as χ_z is not too large.

In other words, we could easily allow the M-type project to also produce cash flow X_z+x, albeit with lower probability. The analysis would remain similar subject to non-essential changes in payoffs.