Capital Commitment

ELISE GOURIER, LUDOVIC PHALIPPOU and MARK M. WESTERFIELD[*]

ABSTRACT

Twelve trillion dollars are allocated to private market funds that require outside investors to commit to transferring capital on demand. We show within a novel dynamic portfolio allocation model that ex-ante commitment has large effects on investors' portfolios and welfare, and we quantify those effects. Investors are under-allocated to private market funds and are willing to pay a larger premium to adjust the quantity committed than to eliminate other frictions, like timing uncertainty and limited tradability. Perhaps counter-intuitively, commitment risk premiums increase with secondary market liquidity and they do not disappear when investments are spread over many funds.

Conflict-of-interest disclosure statement

I. Introduction

Institutional investors' exposure to private market funds amounts to over $12 trillion.[1] These funds span a wide range of investments from real estate to leveraged buyouts, private debt and venture capital. A defining feature of private market funds, irrespective of their focus, is that they require investors to commit capital to fund managers before it is used ("called"), and thus to relinquish control over their portfolio allocation. Capital commitments are large, having trebled since 2008 to a total of $3.2 trillion, and the average delay between commitments and calls is significant at about three years. This paper quantifies the effect of these ex-ante capital commitments on portfolio allocation decisions and investors' welfare.

We solve the dynamic portfolio optimization problem of a risk-averse investor with an infinite horizon and access to stocks, bonds, and private equity funds. We model investments in Private Equity (PE) from the investor's perspective, taking as given the features and other key institutional details of PE contracts. At time 0, the investor commits a positive amount to a PE fund; they do not know when the capital will be called or when investment proceeds will be distributed. We use Poisson processes to model the stochastic timing of capital calls and distributions. The first jump triggers the capital call, and the investor transfers the committed amount to the fund manager or defaults on their commitment. If the investor makes the transfer, the capital is invested by the fund manager. The second jump of the Poisson process marks the time at which the fund manager distributes the proceeds from the fund back to the investor. Then, a new capital commitment can be made to a new PE fund, and the process is repeated.

Our model incorporates both strategic default – the investor can skip a capital call at the cost of lost future opportunities – and access to a secondary market – the investor can sell the claim on their invested capital. We conduct

---

[1] As is customary, the $12 trillion figure represents the sum of the Net Asset Value of all existing funds and of all committed but uncalled capital ("dry powder"). This information is from the Preqin Pro website, as of November 2022.

a unique and thorough calibration of our model. We jointly estimate the set of parameters that best capture the empirically observed speed of capital calls and the cross-sectional distribution of fund performance as measured by the Kaplan-Schoar Public Market Equivalent (PME). We believe this is the first time a comprehensive structural model of PE investments is brought to the data, delivering a quantitative and empirically implementable portfolio choice approach to evaluate these investments. We show that our relatively stylized model closely and simultaneously matches the whole empirical distribution of both capital calls and fund performance.

Our setup allows us to define and decompose the liquidity frictions related to commitment. The delay between a capital commitment and call is stochastic, and we call the associated risk *commitment-timing* risk. Next, because of public market movements while waiting for the capital call, the quantity of capital committed *as a fraction of wealth* is stochastic;[2] we label the risk associated with this friction *commitment-quantity* risk. For both risks, we define i) the welfare cost as the one-off amount of wealth the investor would give up to remove the risk from the economy; and ii) the return premium as the permanent PE return loss the investor would accept to remove the risk.

Our central result is that the cost of *commitment-quantity* risk is large while the cost of *commitment-timing* risk is negligible. To switch to an economy without *commitment-quantity* risk, i.e. one in which they would be able to adjust their PE allocation at the time of capital call, the investor is willing to pay 1.25% of their initial wealth (i.e., 24% of their optimal PE commitment), or accept a permanent loss of 1.10% of PE returns. This cost is driven by the fear of moving away from the target PE allocation. A PE allocation that is too large relative to liquid wealth leads to a reduction in consumption relative to liquid wealth because the investor cannot consume out of their PE stakes during the holding period. During the commitment

---

[2]During the time period between commitment and capital call, investors face stock volatility, causing the fraction of wealth committed to be sub-optimal at the time of the capital call. For example, following a decline in stock prices during the commitment period, the amount called is larger than the optimal amount.

period, the investor looks forward and anticipates the welfare loss incurred by a sub-optimal PE allocation after the capital call. Consumption volatility is further impacted by management fees during the commitment period, which are proportional to commitment and paid out of liquid wealth. To avoid the possibility of the PE allocation becoming too large, the investor under-commits to PE: on average, they would like to nearly double their PE allocation at the time of capital call.

Our results offer a novel rationalization for the increased offering of co-investment opportunities to investors by fund managers at the time of capital call. Co-investment opportunities are valuable options to increase PE exposures post commitment, whereas the literature mostly presents them as tools to reduce fees.

In contrast to commitment-quantity risk, *commitment-timing* risk carries a cost close to zero. In fact, this cost can be negative for large values of the subjective discount factor. This result is surprising because timing risk increases the dispersion of PE allocations and therefore amplifies commitment-quantity risk. We explain this result as follows: The investor's utility increases at the time of capital call, and exponential discounting is a convex function of time, making a deterministic time of capital call less valuable than a stochastic time.

Neither the option to strategically default nor the secondary market alleviate commitment-quantity risk. Given the investor's under allocation to PE, it is almost never optimal for the investor to surrender future opportunities by strategically defaulting on a capital commitment. Similarly, the secondary market, as a tool to liquidate PE positions with a haircut, is relatively unimportant to the investor. As the investor cannot sell partial stakes in a fund[3], they rarely find themselves with such excess holdings that they would like to sell. When we extend our model to an infinite number of funds, partial sales are de-facto allowed. However, the investor still uses the secondary market rarely. When they are over-allocated to PE, they are better

---

[3]We discuss how our assumptions on the private equity secondary market match practice in Section E.

3

off stopping their commitments to new funds, decreasing the PE allocation through calls of existing commitments.

Our model further implies that the development of a PE secondary market increases the investor's willingness to pay to alleviate commitment-quantity risk. When the secondary market is more liquid, there is a small welfare gain and an increase in PE allocation, which in turn increases the welfare cost and return premiums associated with commitment-quantity risk. The liquidity of the secondary market and commitment risk are thus complements.

Increasing the number of funds allows us to study the effects of diversification. Our calibration of the one-fund model leads to an optimal PE commitment of 5.2% of wealth. This optimal commitment increases to 8.5% with two funds. With an infinity of PE funds, at the steady state, 21.9% of wealth is allocated to PE.[4] However, the effect on risk premiums is negligible. The investor accepts a permanent PE return reduction of 0.86% to go from one PE fund to two PE funds. This is *less* than the premium associated with commitment-quantity risk. In addition, going from one fund to two funds hardly decreases the return premium of commitment-quantity risk, from 1.10% to 0.79%, and access to an infinity of funds only brings it down to 0.74%.

Although the investor is able to smooth both cash flow shocks and investment timing – PE cash flow risk and investment timing risk are idiosyncratic – two effects hinder diversification. First, commitment-quantity risk is driven by the denominator of the commitment-to-wealth ratio. Commitments are constant, but the investor's liquid wealth is volatile, and the denominator is the same for all funds. Thus, even with an infinity of funds, the investor cannot diversify commitment-quantity risk away. Second, investing in multiple funds creates the potential for a *funding mismatch*. Increasing investment

---

[4]Aggregate asset allocation across endowments and foundations as of March 2020 according to the Bank of New York Mellon Corporation was 31.1% in listed equity, 16% in fixed income, 18.8% in hedge funds, and 16.9% in private equity. From https://www.pionline.com/interactive/larger-endowments-foundations-lean-private-equity-allocations

across several funds means using distributions from earlier funds to meet later capital calls. However, if one fund distributes late and another calls early, the investor may be short of liquid assets. Thus, capital commitment remains relevant, even when the investor has access to multiple funds.

We further extend the model to allow for liquidity cycles in which private and public return moments co-vary with call and distribution intensities. Adding liquidity cycles has a large effect on welfare, a smaller effect on portfolio allocation, and an even smaller effect on commitment risk premiums. In our calibration, the bad state features lower returns and higher volatility for both public and private equity, higher correlation between public and private equity, longer average commitment and holding periods, and a larger haircut on the secondary market. In this setup, the return premium remains similar to the no-cycle premium in both states. However, if private and public equity returns are no longer both low in the same liquidity state, any switch in the liquidity state increases the investor's desire to adjust their PE exposure, which exacerbates commitment-quantity risk. In this case, the return premium of commitment-quantity risk slightly increases in both states.

We build on a literature that studies the optimal portfolio choice problem in the presence of illiquid assets. Illiquidity is often defined as the inability to trade an asset during a given period of time, e.g., Longstaff (2001); Kahl, Liu, and Longstaff (2003); Longstaff (2009); Gârleanu (2009); Dai, Li, Liu, and Wang (2015). Recent papers model the specific illiquidity features of private equity funds. In Sorensen, Wang, and Yang (2014), a (single) private equity fund is acquired at time 0, hence the capital is immediately invested, but this investment cannot be traded. The fund is liquidated at maturity T, which is finite and known ex ante. In Ang, Papanikolaou, and Westerfield (2014), an illiquid asset cannot be traded during stochastic periods of time. They illustrate how trading illiquidity can create funding illiquidity, and the resulting portfolio effects and welfare costs are found to be large. Dimmock, Wang, and Yang (2023) allow the agent to liquidate their positions in the illiquid asset on a secondary market at a cost and evaluate the "endowment model" used by some institutions that invest in alternative assets. Bollen and Sensoy (2021) extend the analysis of Sorensen, Wang, and Yang (2014) by

allowing for a secondary market for partnership interests. These papers focus on illiquidity by constraining an investor to hold an illiquid asset over a period of time, which is either deterministic or stochastic.[5] Capital committed to the illiquid asset is immediately invested. Thus, the central feature of private market funds, ex ante capital commitment, is not modelled.

In a contemporaneous paper, Giommetti and Sorensen (2020) model a private equity portfolio in which capital is gradually called and distributed from a composite private equity fund. Capital commitments are therefore implicitly embedded in their model. Their central finding is that the optimal allocation to private equity is not sensitive to risk aversion due to the nature of private equity funds' illiquidity. They also find that this result depends on the liquidity of the secondary market.

II. Institutional Setup

A. Investment Vehicles

Private market investing spans the following investment strategies: Leveraged Buy-Out, Venture Capital, Growth Equity, Private Debt, and Real Assets (real estate, infrastructure, timber, natural resources). In this subsection, we detail the three broad routes that institutional investors have to invest in private markets.[6]

A.1. Blind-Pool Funds

Most of the investments in private markets are made via finite-life closed-ended blind pools of capital, which are structured as private limited partnerships and simply referred to as funds. A company (e.g., KKR & Co. Inc.) acts as the General Partner (GP) for the fund (e.g., KKR XII), and capital

---

[5]Korteweg (2019) and Korteweg and Westerfield (2022) survey this literature.

[6]See also Korteweg and Westerfield (2022) and Phalippou (2021) for a survey of the institutional details associated with private market fund investor issues and the associated academic literature.

is provided by the Limited Partners (LPs). LP interests in a fund cannot be traded; but they can be transferred to another investor with the consent of the GP. The commitment is 'blind' in the sense that LPs do not have a say about whether an investment should be made or not. A fund is a pool of ten to twenty investments.

During a fund-raising period that spans three to eighteen months, a GP seeks capital for its fund. LPs bear a significant due diligence cost to decide whether or not they commit capital (Da Rin and Phalippou (2017)); if they do, they agree to provide cash on demand to the fund, up to their committed amount during a pre-specified "investment period." When the GP ends its fundraising, it has its "final close", and the year this occurs is called the fund vintage year. The time between capital commitment and deployment ("capital calls") is long and spans both the fundraising period and the investment period. This investment approach is sometimes called Commitment-and-Drawdown.

GPs are highly specialized agents who devise a value-add plan for each investment. They are said to pursue a 'buy-to-sell' strategy, i.e. their main objective is to increase the asset value and sell as soon as this value-add can be cashed in. LPs have no say on the timing of asset sales, just as they have no say on the timing of capital calls.

A.2. Solo Investing

Solo investments – as coined by Fang, Ivashina, and Lerner (2015) – are 'direct' ownership stakes taken by asset owners into companies. This route is common for assets that do not require a value-add plan (so-called 'core' assets); for example, New Hampshire's Great North Woods (Yale Endowment), London O2 arena (Trinity College Cambridge). However, Fang, Ivashina, and Lerner (2015) also identify some solo investments among the LBO and VC investments of the large institutional investors in their dataset.

With solo investments, institutional investors target an amount of capital ex-ante, search for an opportunity, and eventually deploy the capital. The time between commitment and deployment varies but it may take a few

months. These investments are typically intermediated by specialized agents (e.g., generalist real estate brokers such as Savills in the UK), do not require a value-add plan and are buy-and-hold investments, not buy-to-sell.

A.3. Discretionary Vehicles

Discretionary investments are buy-to-sell investments that are proposed separately by intermediaries to prospective asset owners. The intermediaries devise and implement a value-add plan, and the prospective owner does costly due diligence each time before deciding to opt in or to pass. There are two main sub-categories of discretionary vehicles:[7]

Pledge funds: Fund participants pledge to contribute capital to a series of investments, but have the right to opt out of specific investments. These structures are observed more in certain regions (e.g. in Asia), for certain types of assets (e.g. real estate, venture capital), and with less established fund managers.

Fund co-investments: Fund co-investment opportunities allow LPs to add capital to a deal when the GP makes a capital call. They are restricted to LPs that are already in the fund; this effectively gives LPs the opportunity to take a greater stake in some of the fund's investments. Pre-2008, co-investment invitations were limited to large LPs; but since then, they are widespread.

---

[7]The label 'discretionary vehicles' was coined by Lerner, Mao, Schoar, and Zhang (2022); they define it as follows: "co-investments into individual companies by one or more LPs; solo investments by LPs in previously private capital-financed companies; pledge fund structures where transactions are funded by the LP on a deal-by-deal basis (sometimes raised by groups that have encountered poor performance and who encountered difficulties raising a traditional fund); co-investment or overage funds that are raised alongside a main fund; and co-sponsored transactions between LPs and GPs. We also include co-investment funds raised by funds-of-funds and other intermediaries." Note that some funds allow investors to add capital regularly over time; they could also be considered discretionary vehicles (e.g. Tiger Global).

B. Why are Blind-Pool Fund Structures Dominant?

The vast majority of private market investments are made via blind-pool funds. In dollar terms, Lerner, Mao, Schoar, and Zhang (2022) report that private market investments are split: 93% in blind-pool funds and 7% in discretionary vehicles; but they exclude real assets, private debt, and solo investments. Solo investments are at best as large as discretionary investments.[8] Blind pools are intermediated and pool capital commitments.

Intermediaries are ubiquitous in financial markets and their existence has been justified by two features. First is transaction cost minimization: pooling capital across multiple agents reduces the per unit cost due to the presence of fixed costs. Second is the information advantages of specialized investors, which is probably significant in private markets. These benefits are counterbalanced by agency frictions.[9] This tradeoff is consistent with the fact that most direct investments are core real asset investments, which are seen as the least complex investments, whereas nearly all LBO investments, which are generally perceived as more complex, are done via a specialized intermediary.

Axelson, Strömberg, and Weisbach (2009) argue that ex-ante capital commitments can be a second best optimal contracting solution. In their model, fund managers have skills in identifying and managing potentially profitable investments, but as they have limited liability, they have an incentive to overstate the quality of potential investments when they raise financing. This agency problem is minimized when capital is committed ex ante to finance a number of future projects rather than when capital is raised on a deal-by-deal basis. This model, therefore, provides a rationale for why investors would

---

[8]In the dataset of Fang, Ivashina, and Lerner (2015), co-investments alone are larger than solo investments, but real assets and private debt are excluded.

[9]GPs deploy capital too quickly at market peaks (Axelson et al., 2013), exit investments prematurely (Barrot, 2017), invest sub-optimally near the end of the investment period (Arcot, Fluck, Gaspar, and Hege (2015), Degeorge, Martin, and Phalippou (2016); and LPs need to provide GPs with "liquidity insurance" in bad times (Lerner and Schoar (2004)).

accept a fund structure despite the cost associated to it.

In addition, in practice, deal-by-deal structures suffer from severe short-comings: "the fund manager will not have existing contractual commitments from investors, that can be called down at very short notice, and this can affect the ability of the manager to commit to underlying transactions in a timely manner. Clearly, entering into a binding underlying purchase contract cannot be finalized until the necessary capital has been raised, as this would give rise to a risk of a breach of contract if the funding cannot ultimately be obtained. If the fund manager is competing against another potential purchaser for an asset, and such other purchaser already has guaranteed funding in place, the vendor may prefer to deal with such other purchaser."[10] As Fang, Ivashina, and Lerner (2015) conclude: "In sum, the different approaches to private equity investing - the traditional intermediated partnership vs. direct investing - present a tradeoff between cost and investment quality."

C.   Capital Calls and Distributions

We now discuss some of the specific institutional features of intermediated blind pools. A capital call is made by the GP on the LPs and is in connection to either a fee payment or an investment. The timing of capital calls is uncertain; LPs only know an ex-ante specified investment period, during which most capital calls should occur. The length of the investment period depends on the investment type. Leveraged buyout funds typically have a five year investment period, with some capital called afterwards for fee payments or follow-on investments in existing portfolio companies. For venture capital funds, the investment period is typically longer to allow for large sums to be invested in later stage rounds for successful portfolio companies. To reduce the frequency of capital calls, GPs often pool some of them and bridge-finance using credit facilities with LP commitments as collateral.

---

[10]Source: https://www.harneys.com/hubs/offshore-funds/the-art-of-the-deal-by-deal/.

The capital distribution period is flexible, spanning the entire life of the fund, including an overlap with the investment period. When an investment is exited, the payout is distributed to LPs and cannot be recycled to make a new investment, but there are some exceptions.

Funds' life is set to ten years but there are multiple circumstances under which funds obtain extensions. Most funds are not fully liquidated by their twelfth year, which shows that there is no hard deadline in practice.[11] There is a wide dispersion in the number of contemporary fund commitments held by institutional investors. At the high end, CalPERS, which is one of the most active PE investors, reports a total of 311 commitments to PE over the last 23 years, i.e., about 65 simultaneously active commitments. At the lower end, many small Endowments and Family Offices only have one or two active commitments. Using Preqin data, we find that the PE allocation as a fraction of AUM is positively correlated with the (log) number of funds (non-tabulated), with a correlation coefficient at 0.24 (t-stat of 17). Thus, institutions differ greatly in their degree of diversification within PE (see also, e.g., Cavagnaro, Sensoy, Wang, and Weisbach (2019)), with many institutions facing lumpy stochastic capital calls.

D. Defaulting on Commitments

The stated penalties for default, as specified in limited partnership agreements, are high (Banal-Estañol, Ippolito, and Vicente (2017)). Penalties include forfeiture of some or all existing investments in the fund, and impossibility to invest in subsequent funds. Perhaps as a result, default is rare. We do not know of any major PE investor that has defaulted on their PE commitment. There is also anecdotal evidence that LPs are willing to take significant and costly actions to avoid default. These costly actions include,

---

[11]In the Preqin dataset of US-focused funds, 74% of the funds were not liquidated after 12 years (73% of buyout funds and 76% of VC funds). Further, Barrot (2017) finds that the type of investment is influenced by the fund age. Earlier investments tend to be in younger companies.

for example, redeeming capital from other investments despite low overall liquidity, selling their fund stakes on the secondary market at large discounts, and issuing high yield bonds.

To illustrate the extent and cost of the default avoidance strategies, here is a typical account of how PE investors fared during the 2008 crisis: "A growing set of limited partners find themselves short on cash amid the financial crisis – and thus are scrambling for ways to make good on undrawn obligations to private equity vehicles. Among those in the same boat: Duke University Management, Stanford Management, University of Chicago and University of Virginia... Brown, whose $2.3 billion endowment has a 15% allocation for private equity products, is apparently thinking about redeeming capital from hedge funds to raise the money it needs to meet upcoming capital calls from private equity firms... Carnegie, a $3.1 billion charitable foundation, is also in a squeeze. Its managers have been calling on commitments faster than expected, while distributions from older funds have slowed down, creating a cash shortfall. As for Duke, the university's endowment has been named as one of the players most likely to default on private equity fund commitments. That partly explains a massive secondary-market offering that the school floated last month, as it sought to raise much-needed cash and get off the hook for undrawn obligations by unloading most of its $2 billion of holdings in the sector... Some of the bigger investors are considering tapping credit facilities to meet near-term capital calls."[12]

E. The Secondary Market

Before 2006, the secondary market for fund stakes was quasi non-existant due to contractual restrictions on transfers (see Lerner and Schoar (2004)). This market then grew quickly from an annual turnover of $10 billion in 2006

---

[12]From the magazine 'Private Equity Insider' in its issue of November 5, 2008. See also Barron's, 6/29/2009,"The Big Squeeze"; Forbes, 10/24/2009, "Did Harvard Sell At the Bottom?"; Institutional Investor, 11/4/2009, "Lessons Learned: Colleges Lose Billions in Endowments."

to over \$100 billion in 2021. Yet, \$100 billion volume still represents less than 1% of the \$12 trillion allocated to private market funds.[13]

Nadauld, Sensoy, Vorkink, and Weisbach (2019) report an average discount to the reported Net Asset Value (NAV) of 13.8% (9% since 2010). In addition, and importantly, they find few transactions of funds during their investment period. Only 9.8% of the transactions occur with funds that are less than three years of age, and even within this category most transactions occur for funds that are three years old. When funds are three years old, they are already nearly two-third invested. It is therefore rare to observe the sale of a pure capital commitment.

Sellers are LPs who transfer their entire stake in a fund: partial sales are rare. Buyers are specialized intermediaries managing dedicated vehicles that are structured as blind-pool funds. These buyers raise equity from asset owners and borrow capital to buy fund stakes on the secondary market. LPs usually buy stakes on the secondary market through these intermediaries rather than directly, with additional charges and delays. Hence, secondary markets mostly allow for downward adjustments in private market allocations. For upward adjustments, investors need to use discretionary or solo investments.

---

[13]For turnover data, see https://www.jefferies.com/CMSFiles/Jefferies.com/Files/IBBlast/Jefferies-Global-Secondary-Market-Review.pdf

III. Model

We model investment portfolios that combine private equity with liquid risky and riskless assets. Our setup is designed to capture important institutional details from Section II, and it allows for one, two, or an infinity of private equity funds.

A. The Liquid Assets

There are two liquid assets in the economy that can be rebalanced continuously at no cost: a risk-free bond, which captures the fixed-income market, and a risky stock, which captures the public equity market.[14]

The price B_t of the bond appreciates at constant rate r,

$$dB_t = rB_tdt, \tag{1}$$

and the stock price P_t follows a geometric Brownian motion,

$$\frac{dP_t}{P_t} = \mu\,dt + \sigma dZ_t^L, \tag{2}$$

where Z_t^L is a standard Brownian motion associated with liquid public markets, μ is the return drift, and σ is the return volatility.

The investor's liquid wealth W_t is the sum of their holdings in the stock and bond.

---

[14]We consider an information structure that obeys standard assumptions. There exists a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ supporting the vector of four independent Brownian motions $Z_t = (Z_t^L, Z_t^{PE}, Z_t^{1\perp}, Z_t^{2\perp})$, and two independent Poisson processes $M_t = (M_t^1, M_t^2)$. Some stochastic processes may be unused, depending on the number of private equity funds. $\mathcal{P}$ is the corresponding measure and $\mathcal{F}$ is a right-continuous increasing filtration generated by $Z \times M$. Following Dybvig and Huang (1988) and Cox and Huang (1989), we restrict the set of admissible strategies to those that satisfy the standard integrability conditions. All policies are appropriately adapted to $\mathcal{F}_t$.

B. Modelling Private Equity

Investors can allocate capital to private equity funds. We begin by describing the model for one fund and then extend it to include multiple funds.

As illustrated in Figure 1, the fund manager collects capital commitments from the investor at time τ_0, calls the committed capital and invests it at time τ_C, and then distributes the value of the investment at time τ_D. We call the *commitment period* the time period $[\tau_0, \tau_C)$ (the fund is in state C), and the *holding period* the time period $[\tau_C, \tau_D)$ (the fund is in state D). After the fund distribution, the investor makes their next commitment and the process repeats to infinity.

Figure 1. Timeline of a fund's life.

At time τ_0, the investor commits a positive amount $X_{\tau_0} \geq 0$ to the private equity fund. This commitment is a promise to make capital available when the manager calls it at time τ_C. The commitment X_{τ_0} cannot be changed after time τ_0, meaning $dX_t = 0$ until the committed capital is called and invested at τ_C. During the commitment period $[\tau_0, \tau_C)$, fees are paid out of the investor's liquid wealth to the fund manager at rate $f X_{\tau_0} dt$, but no investment is made.

We use a Poisson process to model the timing of capital transfers between the investor and the fund manager. The process has intensity λ_C during the commitment period. A jump triggers the capital call and the end of the commitment period, at which time the investor transfers X_{τ_0} of liquid wealth to the fund manager.

In a slight abuse of notation, we denote by X_t the amount of capital committed between times τ_0 and τ_C, and we use X_t again to refer to the

15

net-of-fee amount of capital invested in the fund after τ_C.

After capital is transferred and invested, the value of the private equity asset, net of all fees, evolves as a geometric Brownian Motion, with drift ν, and volatility ψ:[15]

$$\frac{dX_t}{X_t} = \nu dt + \psi dZ_t^X, \tag{3}$$

where $dZ_t^X = \rho_L dZ_t^L + \sqrt{1 - \rho_L^2} dZ_t^{1\perp}$ and $Z^{1\perp}$ is the idiosyncratic shock associated with the fund. This specification implies that the correlation between public and private equity is ρ_L, and the beta of a private equity fund is

$$\beta = \rho_L \frac{\psi}{\sigma}. \tag{4}$$

We use the Poisson process again to model the timing of capital distributions. During the holding period $[\tau_C, \tau_D)$, the intensity of the Poisson process is λ_D, and a jump triggers capital distribution. The private equity investment is fully exited, and the investor receives the value of the fund, X_{τ_D-}. The Poisson process resets, and the investor is immediately able to make a new capital commitment to a new private equity fund[16], i.e. $\tau_D^{fund\ i} = \tau_0^{fund\ i+1}$.

In our setup, the uncertainty around capital calls and distributions is modelled with two random times for the private equity fund, τ_C and τ_D. These two random times represent two sources of market incompleteness. Even if the liquid asset and the private equity fund had fully correlated returns, or if the investor had access to the derivatives market, the investor would not be able to hedge the risk coming from the random times, and the market would still be incomplete.

Our model with one fund relies on two assumptions that ensure analytical tractability. First, there is a single capital call equal to the committed

---

[15]To keep the model parsimonious, during the holding period we do not model the management fee and carried interest separately from returns. Instead, we assume that the net-of-fees value X_t follows a geometric Brownian motion.

[16]Without a pledge to private equity, the investor's opportunities are constant, so it is never optimal for the investor to wait to commit.

amount, as opposed to having capital calls spread across the investment period. Second, there is a single payout.

Below, we generalize this basic model in two important directions. First, we allow for multiple funds, including an infinite-fund limit. Second, we allow for private equity cycles in which parameters, including returns and waiting times, are allowed to vary over time. In all cases, our representation allows us to use numerical methods based on Markov chain approximations to solve the ODEs and PDEs associated with the portfolio allocation problem.

Our model setup is flexible enough to allow for the existence of a secondary market. We assume that during the holding period, the investor can sell their invested private equity on a secondary market, receiving αX_t. $0 \leq \alpha \leq 1$, and $1 - \alpha$ is the haircut. After the sale, the investor waits until the end of the fund's life τ_D, then makes a new commitment, starting the process over.

The investor can strategically default on their capital commitment either at the time of the call or any time before. The consequences of default are that the investor does not turn over the capital and stops paying the associated fee, but the investor is banned from accessing private equity in the future – which is a realistic feature (see Section D).

C.   The Investor's Problem

The investor continuously rebalances their liquid wealth between the two liquid assets, and consumes out of liquid wealth at rate $c_t = C_t/W_t$. We denote with θ_t the fraction of liquid wealth allocated to stocks, so the evolution of the investor's liquid wealth is given by:

$$\frac{dW_t}{W_t} = \left(r + \left(\mu - r\right)\theta_t - c_t\right)dt - \mathbb{1}_{S=C}f\frac{X_{\tau_0}}{W_t}dt + \theta_t\sigma dZ_t^L - \frac{dI_t}{W_t} \qquad (5)$$

where dI_t denotes any transfer between liquid wealth and illiquid wealth. Throughout the paper we will use $\mathbb{1}$ as an indicator variable. Thus $\mathbb{1}_{S=C}f\frac{X_{\tau_0}}{W_t}$ denotes fees that are paid during the commitment period.

17

The value function is given by

$$F(W_t, X_t, S_t) = \max_{\{\theta, X, c\}} \mathrm{E}_t \left[\int_t^\infty e^{-\delta(u-t)} U(C_u) du \right], \tag{6}$$

subject to (3) and (5). We use δ to denote the subjective discount factor and $S_t = \{C, D\}$ to denote the state. The investor has a standard power utility, i.e., $U(C) = \frac{C^{1-\gamma}}{1-\gamma}$, with $\gamma > 1$. Given our assumptions, $W > 0$ and $X \geq 0$ a.s.

At $t = \tau_C$, if the investor has not defaulted on their commitment, we have $dI_t = dI_{\tau_C} = X_{\tau_0}$; i.e. committed capital is called, and X_{τ_0} is transferred out of liquid wealth to private equity. The state changes from $S = C$ (commitment period) to $S = D$ (holding period), and the value function jumps discretely from $F(W, X, S = C)$ to $F(W - X, X, S = D)$.

The investor strategically defaults if the welfare value of the standard Merton problem, $F^{Merton}(W)$, – the solution to our model with access to the liquid stock and bond but without the PE fund– exceeds the continuation value with PE. The investor defaults before a capital call if $F(W, X, S = C) < F^{Merton}(W)$ and upon a capital call (at τ_C) if $F(W - X, X, S = D) < F^{Merton}(W)$.

During the holding period, the investor can sell their stakes on the secondary market, and they do so if $F(W, X, S = D) < F(W + \alpha X, 0, S = D)$.

When $t = \tau_D$ is reached, we have $dI_{\tau_D} = -X_{\tau_D-}$; i.e. capital is paid out and X_{τ_D} is transferred from PE to liquid wealth. Then, the investor chooses their level of committed capital to the new fund and the state changes from D to C. The value function jumps discretely from $F(W, X, S = D)$ to $\max_{X'} F(W + X, X', S = C)$. At all other times, $dI_t = 0$.

The investor's value function, optimal consumption and allocation solve the Hamilton-Jacobi-Bellman (HJB) equation given in Appendix IA.A.1. Because the utility function is homothetic and the return processes have constant moments, the value function F is homogenous of degree $1 - \gamma$ in total wealth. We use V to denote total wealth and $\xi \geq 0$ to denote the fraction of

total wealth either committed or invested, so that

$$V = W + X \mathbb{1}_{S=D} \tag{7}$$

$$\xi = \frac{X}{V}. \tag{8}$$

Thus, the investor's value function can be written as the product of a power function of total wealth and a function of the wealth composition:

$$F(W, X, S) = V^{1-\gamma} H(\xi, S). \tag{9}$$

The optimal commitment is given by the following Proposition:

PROPOSITION 1: *The investor's value function can be written as in (9), where $H(\xi, S)$ exists and is finite, continuous, and twice differentiable. Whenever the investor can commit capital, they select $\xi^* \equiv \arg\max_\xi H(\xi, S = C)$, which exists.*

The function H is characterized by the set of ODEs shown in Online Appendix IA.A.1, and our method for generating numerical results is detailed in Online Appendix IA.B.

D. The Illiquidity Frictions

The model presented in Sections B and C defines our **baseline economy (Economy 0)**, i.e. an economy where all liquidity frictions are present. We now define five other economies, each of which corresponds to a situation where one or more of the frictions of our baseline model are modified. These modifications allow us to assess theoretical counter-factuals and isolate the impact of the various private equity investment frictions. The ODEs and PDEs that characterize the solutions to the investor's problem in these five economies are given in the Online Appendices IA.A.2 to IA.A.6.

Economy 1: Deterministic call time, choose quantity when committing.

The first friction is that the time of capital call is unknown. We turn off this *commitment-timing risk* by making the call time deterministic. The

agent commits capital at τ_0, but instead of waiting for a random delay of expected length $\frac{1}{\lambda_C}$, capital is called deterministically at $\tau_C = \tau_0 + \frac{1}{\lambda_C}$. Thus we maintain the average delay but remove the uncertainty about this delay.

Economy 2: Stochastic call time, choose quantity when called

Next, we turn off *commitment-quantity risk*. In the baseline economy, the committed capital is fixed at τ_0, but liquid wealth evolves randomly before the committed capital is called. Thus, *the relative size of the commitment changes*. Investors do not know the fraction of their wealth that they will have to pay out when capital is called. Similarly, the relative size of fees changes because they are assessed out of liquid wealth in proportion to committed capital. To turn this commitment-quantity risk off, we let the investor choose the quantity invested in private equity when the capital is called, instead of at commitment time. So, the commitment-to-wealth ratio ξ^* is chosen at call time τ_C rather than commitment time τ_0. This change removes any risk of default, but the timing of the capital call remains stochastic.

Economy 3: Deterministic call time, choose quantity when called

If neither commitment-timing nor commitment-quantity risk is present, there is a commitment delay but *no commitment risk*.

Economy 4: Immediate private equity access

Absent commitment risk, our model still features a *commitment delay*. The investor needs to wait until call time to access the private equity returns: their active investment time is the holding period of the fund. This restriction can be lifted and the commitment period time brought to zero. Hence, at time τ_D, the investor freely allocates capital between bonds, stocks and private equity and directly enters the fund's holding period.

Economy 5: Deterministic payout time

Economies 1-4 constitute a peeling back of the institutional details associated with commitment risk. We remove the commitment-timing and quantity risks separately and then together, and then we remove the commitment delay as well. For comparison, we also consider making the stochastic distribution time deterministic. In this economy, the investor's capital is called at random time τ_C, but the holding period duration is deterministic, with

20

$\tau_D = \tau_C + \frac{1}{\lambda_D}$. Thus we maintain the average holding period but remove the risk. This economy still contains commitment risk.

E. Economies with several PE funds

We can extend the baseline model presented in Sections B and C to include more than one PE fund. In this sub-section, we describe the economies with two and an infinity of PE funds. Analyzing economies with multiple PE funds is important for at least four reasons.

First, with multiple PE funds, LPs can diversify cash flow risk. Each fund has idiosyncratic risk, and a standard diversification intuition indicates than an investor would prefer to spread their wealth across multiple assets. Second, investors reduce the lumpiness of capital calls and distributions. However, any shock to public equity is a shock to all commitment-to-wealth ratios through the denominator. So while diversifying across funds allows for multiple smaller commitments and smoother call timing, the total outstanding commitment-to-wealth ratio remains volatile. Third, investors can potentially use distributions from earlier investments to fund later investments. Doing so smoothes out their quantity invested and increases the active investment time – the time during which capital is invested and earning returns in PE. However, there can also be a funding mismatch: the risk that earlier distributions will be late or insufficient to fund capital calls. Fourth, in contrast to the one-fund case, investors can partially sell their PE holdings on the secondary market.

E.1. Liquidity Diversification: Investing in Two Funds

As in the one-fund case, the investor chooses an optimal commitment X^i to fund i after this fund distributes its previous round of capital. If the investor defaults on one fund's commitment, we assume that they also exit their second fund, either by defaulting on their commitment or by selling their holding on the secondary market. In both cases, they lose access to PE and their investment opportunity set reduces to the liquid stock and bond, as in the one-fund model.

During each fund's holding period, returns follow (3), with the same expected return ν and volatility ψ. We assume that the Brownian motions that drive fund returns have correlation ρ_L with public market equities and correlation $\rho_{PE} > \rho_L$ with each other. Thus the returns of each fund i during its holding period are given by

$$\frac{dX_t^i}{X_t^i} = \nu dt + \psi dZ_t^i, \tag{10}$$

where $dZ_t^i = \rho_L dZ_t^L + \sqrt{\rho_{PE}^2 - \rho_L^2} dZ_t^{PE} + \sqrt{1 - \rho_{PE}^2} dZ_t^{i\perp}$, Z^L is the public market shock, Z^{PE} is a common PE shock, and $Z^{i\perp}$ is the idiosyncratic shock associated with fund i.

Even if the two funds are synchronized at some point in time, they will rapidly de-synchronize because of the stochastic call and distribution timing. One should thus think of the steady state in this economy as one in which calls and distributions randomly overlap one another. The solution to this problem is given in Online Appendix IA.A.7.

E.2. Limiting Case: Investing in an Infinity of Funds

Appealing to the law of large numbers, we assume that when there are an infinity of PE funds, PE funds make calls and distributions continuously, and the investor makes commitments continuously as well. Since individual funds have commitment periods that are exponentially distributed with parameter λ_C, a fraction $\lambda_C dt$ of funds call capital over the interval dt. Similarly, a fraction $\lambda_D dt$ of funds make distributions over dt.

These assumptions imply that commitment-timing risk does not exist with an infinity of funds. However, commitment-quantity risk remains: the investor's commitments are called over time, and liquid wealth is fluctuating randomly.

Our two state variables are the aggregate capital committed to all PE funds (X_t^∞), and the aggregate invested amount (Y_t^∞). We label the investor's new commitments as $dJ_t \geq 0$. Then, extending (5) and (10), we

have the following dynamics:

$$\frac{dX_t^\infty}{X_t^\infty} = \frac{dJ_t}{X_t^\infty} - \lambda_C dt \tag{11}$$

$$\frac{dY_t^\infty}{Y_t^\infty} = \lambda_C \frac{X_t^\infty}{Y_t^\infty} dt - \lambda_D dt + \nu dt + \psi^\infty dZ_t^\infty \tag{12}$$

$$\frac{dW_t}{W_t} = \left(r + (\mu - r)\,\theta_t - c_t\right) dt + \theta_t \sigma dZ_t^L - f\frac{X_t^\infty}{W_t} dt - \lambda_C \frac{X_t^\infty}{W_t} dt + \lambda_D \frac{Y_t^\infty}{W_t} dt, \tag{13}$$

where $dZ_t^\infty = \rho_L^\infty dZ_t^L + \sqrt{1 - \rho_L^{\infty 2}} dZ_t^{PE}$.

The parameters driving equation (12) are those of an equally weighted portfolio of PE funds, taking the limit as the number of funds in the portfolio goes to infinity. As with the two-fund case, we assume that the shocks of each PE fund have correlation ρ_L with public markets and ρ_{PE} with each other. Then we can calculate the volatility ψ^∞ and correlation with the stock market ρ_L^∞ of an equally weighted portfolio analytically:

$$\psi^\infty = \psi\sqrt{\rho_{PE}} \quad ; \quad \rho_L^\infty = \frac{\rho_L}{\sqrt{\rho_{PE}}}. \tag{14}$$

The investor maximizes their expected discounted utility as in (6), subject to the budget constraints (11)-(13) and the constraint that J_t is non-decreasing.

Next, we define the ratio of committed wealth to total wealth, and the ratio of invested illiquid wealth to total wealth to be

$$\pi_t \equiv \frac{X_t^\infty}{Y_t^\infty + W_t} \quad , \quad \xi_t \equiv \frac{Y_t^\infty}{Y_t^\infty + W_t} \tag{15}$$

As in the case of a finite number of funds, the investor's value function can be decomposed into the effect of total wealth and the effect of wealth composition on the continuation utility:

$$F^\infty(W, X^\infty, Y^\infty) = (W + Y^\infty)^{1-\gamma} H^\infty(\pi, \xi), \tag{16}$$

where the function $H^\infty(\pi, \xi)$ satisfies PDEs given in Online Appendix IA.A.8.

In contrast with the one-fund case, there is an optimal PE commitment for each level of PE investment.

Furthermore, the investor can default at any time, in which case they sell their aggregate investment on the secondary market and lose access to PE. They can also sell any fraction ω of their invested capital at any time on the secondary market, which they do if $F^\infty(W, X^\infty, Y^\infty) < F^\infty(W + \alpha\omega Y^\infty, X^\infty, Y^\infty(1-\omega))$. This corresponds to selling complete positions in some subset of the infinity of funds.

We conduct our welfare analysis at the steady state of the aggregate investment Y_t^∞, i.e., when $E[dY_t^\infty] = 0$.[17] We denote this steady state investment by Y^∞. Using the laws of motion, the steady state aggregate commitment X^∞ is linked to the investment Y^∞ by:

$$\frac{X^\infty}{Y^\infty} = \frac{\lambda_D - \nu}{\lambda_C}, \tag{17}$$

Thus, the steady state ratio of committed to allocated capital is a simple function of the rates of calls and distributions and the mean PE return.

Because commitment-timing risk is eliminated with an infinity of funds, our comparison of illiquidity frictions (from Section D) collapses to the comparison between the baseline economy and Economy 2 in which the investor can choose their commitment at the time of a capital call. With an infinity of funds, Economy 2 allows an investor to immediately add to their PE assets (Y_t^∞) to reach the optimal level of invested capital. Equation (12) becomes:

$$\frac{dY_t^\infty}{Y_t^\infty} = \frac{dJ_t}{Y_t^\infty} - \lambda_D dt + \nu dt + \psi^\infty dZ_t^\infty \tag{18}$$

Because negative calls are not allowed, an investor above their optimal allocation must either wait for distributions for their invested capital to decline, or

---

[17]One might also be interested in the steady state for the Y^∞/W ratio: $E[d\ln(Y_t^\infty/W_t)] = 0$. This change implies using a different location in $\{X, Y, W\}$ space to do the welfare analysis. The overall allocation to PE is slightly higher, with more capital committed and less invested, and the welfare and return premiums are almost the same.

sell a fraction of their invested capital on the secondary market at a discount.

F.   Measuring the Cost of Illiquidity

We define two measures to quantify the costs of the different liquidity frictions described in Section D. First, the **welfare cost to the investor** of any economy A with respect to any economy B, denoted $\zeta^{A,B}$, is the fraction of wealth the investor would be willing to pay at the time of commitment to switch from economy A to economy B while simultaneously adjusting their capital commitment. If $\zeta^{A,B} = 0$, then the investor is indifferent between the two economies. $\zeta^{A,B}$ is the solution to the following equation:

$$H^A(\xi^{A*}, S) = (1 - \zeta^{A,B})^{1-\gamma} H^B(\xi^{B*}, S), \tag{19}$$

where ξ^{A*} and ξ^{B*} denote, respectively, the optimal commitments in economies A and B. We evaluate welfare at the time the agent chooses the allocation, i.e., at time τ_0 for Economies 0, 1 and 5 and τ_C for Economies 2, 3 and 4.

Second, the **return premium** of any economy A with respect to economy B is the additional return of the PE funds that would be needed in economy A to make the investor indifferent between the two economies. The return premium applies to all PE funds, both current and future. If the investor is indifferent between economy B with PE expected returns ν and economy A with expected returns $\nu + \epsilon_{AB}$, then the return premium is ϵ_{AB}.

In the two-fund model, welfare costs and return premiums are computed from the value function evaluated at $\xi^1 = \xi^{1*}$ and $\xi^2 = \xi^{2*}$, i.e. assuming that commitments to both funds are optimal.

The two welfare measures should be interpreted differently. The welfare cost is a one-time payment to switch economies, so it is strongly increasing in the optimal PE allocation. In contrast, the return premium impacts the investor in proportion to the amount allocated to PE, so it is much closer to a *per unit* cost of illiquidity.

IV. Model Calibration & Portfolio Allocation

In this section we provide a detailed calibration of private equity return dynamics. We calibrate for economies with one-, two-, and an infinite- number of PE funds and provide the resulting portfolio and consumption policies in the baseline economy.

A. Model Calibration

We use the past thirty years of data to calibrate our model (1991-2020). The average 3-month Treasury Bill is $r = 0.03$. The mean and volatility of the S&P 500 index log returns at monthly frequency are $\mu = 0.08$, and $\sigma = 0.15$. We use standard values for the investor risk aversion and discount factor: $\gamma = 4$ and $\delta = 0.05$. The discount for PE-funds secondary market sales is set to the average reported in Nadauld, Sensoy, Vorkink, and Weisbach (2019): 13.8%. Management fees during the commitment period are set to $f = 2\%$ of the committed amount (see Metrick and Yasuda (2010)). Our calibration of the private equity return dynamics uses the Preqin dataset with fund cash flows as of the end of year 2020. We select all US-focused private equity funds (venture capital, growth equity, leveraged buyout) raised between 1991 and 2015 (so that they have at least five years of investment activity).[18]

We construct two cumulative distribution functions for fund cash inflows: the empirical distribution and the model-implied distribution. The former is derived directly from the Preqin dataset. The latter is calculated analytically and verified with simulations.

In our model, PE-fund cash *inflows* consist of the management fees during the commitment period plus the investment capital call at time τ_C. Assume

---

[18]We select funds with a size of at least $10 million, at least two capital calls, and at least two capital distributions. The resulting sample contains 1398 funds. Note that Preqin records all cash inflows into any given private equity fund, without a distinction between fee payments and capital invested.

that \$1 is committed to each one of the N funds. The delay from τ_0 to τ_C has an exponential distribution. Therefore, at any time t, the proportion of funds across simulations that have not called yet is $e^{-\lambda_C t}$. The total fees paid by these funds is $N f e^{-\lambda_C t}$. Imposing the law of large numbers, the cumulative cash inflow across is the the sum of the cumulative fees paid by each fund until its capital call, $\int_{\tau_0}^{t} (Nf)(e^{-\lambda_C t}) dt = \frac{Nf}{\lambda_C}\left(1 - e^{-\lambda_C t}\right)$, plus the amount of capital already called, $(1 - e^{-\lambda_C t})N$. As $t \to \infty$, i.e. after all funds have exited their investments, the total cash inflows approach $N\left(\frac{f}{\lambda_C} + 1\right)$. Thus, the cumulative cash inflow at time t, as a fraction of the total, is

$$\frac{\frac{Nf}{\lambda_C}\left(1 - e^{-\lambda_C t}\right) + (1 - e^{-\lambda_C t})N}{N\left(\frac{f}{\lambda_C} + 1\right)} = 1 - e^{-\lambda_C t},$$

which implies an exponential distribution with parameter λ_C.

We search for the λ_C that minimizes the least-square distance between the model-implied and empirical cumulative distributions. The best fit is obtained for $\lambda_C = 0.344$, which corresponds to an average commitment period of about three years. Panel A of Figure 1 shows that the two cumulative distribution functions are very close to one another, with an RMSE of 1.9×10^{-2}, which validates our modeling choice.

To calibrate cash *outflows*, we do not directly use the fund distributions observed in Preqin, as we do for the inflows. The reason is that it takes about fifteen years to observe the complete time series of fund distributions, and we would then be restricted to using a sample of funds raised before 2005. Instead, we use the same sample as above – funds raised up to 2015 – and match the distribution of their performance as of end of year 2020. To measure fund performance, we adopt the most common measure: the Kaplan-Schoar Public Market Equivalent (PME). That is, for each fund, we compute the present value of cash inflows and cash outflows, each discounted using the realized S&P 500 index returns, and we value unexited investments at their reported Net Asset Value (NAV). As we are interested in the distribution of fund-level PMEs, we assign equal weight to all funds.

The model-implied PMEs are obtained by simulating the cash flows of

100,000 private equity funds using i) the return dynamics of equation (3), and ii) draws from Poisson distributions to trigger capital calls and distributions. There are four free parameters in our model: PE expected return (ν), PE volatility (ψ), the intensity of capital distributions (λ_D), and the correlation between private and public equity (ρ_L).

We choose the parameters ($\nu, \psi, \lambda_D, \rho_L$) which minimize the least-square distance between the model-implied and the empirical cumulative distributions. The empirical and model-based cumulative distributions of PMEs are shown in Panel B of Figure 1. The two curves are remarkably close, with an RMSE of 3×10^{-2} for PME $\in [0.5, 2.5]$, and the best fit is obtained with a combination of a relatively high ν and ψ:

$$\nu = 14\%; \quad \psi = 33.5\%; \quad \lambda_D = 0.174; \quad \rho_L = 0.66$$

Our calibrated parameters produce an implied private equity β of 1.47:

$$\beta = \rho_L \frac{\psi}{\sigma} = 0.66 \left(\frac{0.335}{0.15} \right) = 1.47.$$

This β is nearly the same as the 1.43 estimate obtained by Ang, Chen, Goetzmann, and Phalippou (2018b), who use a completely different methodology: Bayesian methods to extract the risk exposures that are most consistent with the observed panels of cash flows. Our calibration implies that the CAPM-alpha of PE is 3.6%, which is close to the asset owner belief of 3.9% expected PE outperformance reported in Ang, Ayala, and Goetzmann (2018a).

Our calibrated λ_D implies an average holding period of about 5.7 years, which is close to the median holding period of 5.3 years that Brown, Harris, Hu, Jenkinson, Kaplan, and Robinson (2021) report.[19]

In the two-fund case, we use the set of parameters described above for the dynamics of each fund, but we also need to calibrate the correlation between the two PE funds, ρ_{PE}. We randomly draw 5000 portfolios of two funds each,

---

[19]They do not report an average, but the skewness of the distribution indicates that it would be higher than the median.

with replacement. We calculate the PME for each of these portfolios, and match the PME distribution to its model-implied counterpart again minimizing the least-squared distance between distributions. We obtain a correlation between PE funds $\rho_{PE} = 0.68$.

We calibrate the infinite-fund problem by using analytic extensions of the parameters above. In the infinite-fund problem, the investor is continuously active and earning PE returns. We assume an equally weighted portfolio and take limits as the number of funds in the portfolio goes to infinity. The volatility and the correlation of the PE portfolio with the stock market can be calculated following equation (14):

$$\psi^\infty = \psi\sqrt{\rho_{PE}} = 0.335 \times \sqrt{0.68} \approx 0.276 \;\; ; \;\; \rho_L^\infty = \frac{\rho_L}{\sqrt{\rho_{PE}}} = \frac{0.66}{\sqrt{0.68}} \approx 0.80. \tag{20}$$

This correlation estimate matches the most commonly used estimate in practice, which is the one given by Blackrock on their capital market assumption webpage.[20] Table 1 summarizes the parameter values that we use. It is remarkable that although we use a parsimonious model, we not only match the distributions of PMEs to their empirical counterparts, but also generate calibrated parameters that are in line with the literature. These results give us additional confidence in our model and the associated counter-factuals.

B.   Portfolio Allocation in the Baseline Economy

B.1. Allocation to Private Equity

At time τ_0, the beginning of each PE fund's life, the investor chooses the optimal commitment to PE as a fraction of total wealth, ξ^*. Following Proposition 1, ξ^* is chosen so that it maximizes the value function that prevails during the commitment period. After commitment, fluctuations in

---

[20]https://www.blackrock.com/institutions/en-us/insights/charts/capital-market-assumptions. The model used by Blackrock to generate this correlation is not publicly available, but they also use PE fund cash flows to infer the correlation.

liquid wealth make the committed amount as a fraction of total wealth, ξ_t, move away from ξ^*. At the time of the capital call and during the holding period, ξ_t denotes the amount invested in PE as a fraction of total wealth. We refer to ξ_t as the PE allocation in both cases.

Figure 2, Panel A shows the portion of the agent's value function related to their investment in PE, $H(\xi, S)$ from equation (9). The dashed (solid) line represents the function during the commitment (holding) period. The investor optimally chooses to commit $\xi^* = 5.2\%$ of wealth. If public markets decline, so that liquid wealth decreases and ξ increases, welfare declines rapidly. The investor strategically defaults on their commitment if the PE allocation reaches 20.8% of wealth during the commitment period (black circle). The investor does not wait for a capital call before defaulting; they avoid paying management fees by defaulting early, but they lose access to PE. If allowed, the investor would pay 1.18% of their total wealth to re-optimize their PE allocation from 20.8% to 5.2% of wealth.

However, Panel B shows that the likelihood of reaching a 20.8% PE allocation is nearly zero, so default hardly ever happens. In fact, the PE allocation does not vary much during the commitment period and remains smaller than 7.4% with a 99% probability. Thus, the states of the economy that lead to strategic default are both rare and important for welfare. This matches the institutional details in Section D.

When capital is called, the value function jumps up to the solid black line (Panel A). At that point in time, the optimal PE allocation is 9.1%, compared to an earlier optimal capital commitment of 5.2% (and a 99th percentile of 7.4%). Thus, the investor chooses an optimal commitment that results in a significant under-allocation to PE. This is despite the fact that liquid wealth drifts up on average, meaning that the investor's commitment as a fraction of wealth declines on average from 5.2% by the time capital is called.

In the holding period, the agent cannot adjust their portfolio nor consume out of their illiquid wealth. The welfare costs to having a sub-optimal allocation to PE are high. With a 13.8% discount (calibrated), the investor sells their stakes on the secondary market only when the PE allocation ξ reaches

58.4% of wealth (black square on Panel A), which is a rare event (Panel B). If allowed, the investor would pay 8.8% of their wealth to re-optimize their PE allocation from 58.4% to 9.1% of wealth.[21] Like strategic default, the states of the economy that lead to secondary market sales are both rare and important for welfare. These features are consistent with the institutional details in Section E: secondary market volume is low relative to aggregate allocations.

B.2. Liquid Portfolio and Consumption Policies

Panel C of Figure 2 shows how the investor alters their consumption policy if they move too far from their optimal portfolio composition. During the commitment period (dashed line), consumption reaches its maximum at the optimal PE commitment. During the holding period (solid line), the investor consumes more than during the commitment period for a given PE allocation, but the consumption rate drops rapidly as the investor approaches the threshold at which they sell on the secondary market.

Panel D describes the liquid asset allocation. The stock allocation fluctuates around 55.5% during the commitment period, i.e., 44.5% of the portfolio is invested in the bond. This large allocation to the bond would be only slightly smaller (44.2%) without PE. Recall (previous section) that strategic default occurs when the PE commitment reaches 20.8% of wealth. Thus, the investor does not allocate more to the liquid bond to avoid default.

However, the option to default creates a near convexity in the value function (Figure 4, Panel A).[22] Thus, the investor takes more risk – tilts their

---

[21]For comparison, liquidating a PE allocation of 58.4% of wealth at a haircut of 13.8% implies a total cost of 8.1% of total wealth. The investor is willing to pay 8.8% to re-optimize – to move to 9.1% instead of zero. The difference, 0.7%, is how much the agent would pay to re-optimize from an initial holding near zero.

[22]Note that while H is convex in ξ, ξ is a composition of liquid and illiquid wealth, and the value function as a whole is not convex in liquid wealth. We discuss the default option further in Section A.

allocation toward the liquid stock and away from the liquid bond – as they approach default (Figure 2, Panel D).

In contrast, during the PE holding period, the investor's allocation to public market equity declines strongly and monotonically with their PE exposure. This is simple hedging: PE and the liquid markets are correlated and an excess allocation to PE is associated with more volatile consumption. So, when public markets decline, the relative allocation to PE increases, and the investor responds by reducing their stock allocation, thereby taking less risk with their liquid assets.

Both consumption and allocation to liquid assets are consistent with private equity changing the concavity of the investor's value function. Capital calls are good news, and consumption is higher during the holding period than the commitment period. However, an unbalanced portfolio reduces the investor's welfare, and consumption is more sensitive to market movements during the holding period. The investor holds a portion of their liquid assets in the bond to finance their commitments, but gambles to avoid default. Once the investor's commitment has been called, they modify their stock holdings to control their overall investment risk exposure.

C. The Illiquidity Stack

In this section, we quantify the impact of the different liquidity frictions on the investor's portfolio allocation and welfare by solving the investor's problem in the five economies described in Section D. The objective is to understand which frictions cause the under-allocation to PE and how costly these frictions are for the investor.

C.1. Commitment-Timing Risk

We start by examining the difference between Economies 0 and 1, i.e., we turn off call timing risk by making the capital call time deterministic instead of stochastic. Results in Table 2 and Figure 3 show that eliminating timing risk has little impact on the optimal commitment, consumption policy, and stock allocation. The welfare cost and return premiums associated with

commitment-timing risk are close to zero. In fact, Figure 10 shows that for larger values of the subjective discount factor δ, the welfare cost becomes negative, implying that the agent *prefers* uncertainty about the timing of capital calls then.

This result is surprising. Indeed, a random call time makes it possible to have either a short or a long commitment period. In the former case, the PE allocation stays close to its optimal level. In the latter case it can significantly depart from this level as the value of the stock changes. In contrast, without timing risk, the commitment period is always the average duration. The distribution of the PE allocation hence depends on the length of the commitment period, and thus on the presence of timing risk. We would expect the agent to prefer certainty over the distribution of the PE allocation because a stochastic capital call time induces uncertainty in the ability of the investor to fund both the capital call and consumption. Because the PE allocation ξ varies over time during the commitment period (due to fluctuations in liquid wealth), the investor is more likely to default on their commitment or reduce consumption if the commitment period is longer.

Two competing forces can explain the sign and magnitude of the commitment-timing risk premium. On the one hand, when capital is called, the value function jumps up, as shown by Panel A of Figure 2. A stochastic capital call time implies that there is uncertainty about the timing of this utility gain. The key fact is that the expected present value of a utility gain is *increasing* in uncertainty over its timing because discounting, $e^{-\delta t}$, is a convex function of time. Jensen's inequality implies $\mathbb{E}\left[e^{-\delta \tau_C} U\right] > e^{-\delta \mathbb{E}[\tau_C]} U$. This feature pulls the cost of commitment timing risk down, and the effect increases in δ.

Our results show that the first force, i.e., the changing distribution of PE allocation, has a small effect and can be outweighed by the trade-off with the convex value of uncertainty. Importantly, whether the net effect is positive or negative, the cost of timing risk remains close to zero, of the order of 10^{-4}.

C.2. Commitment-Quantity Risk

Next, we compare Economies 0 and 2, i.e., we enable the investor to adjust their committed amount (upward or downward) upon capital call. This comparison allows us to evaluate the impact of commitment-quantity risk. Results in Table 2 show that switching off quantity risk leads to a significant increase of the optimal PE commitment. Instead of a commitment of 5.2% at time τ_0, the investor optimally commits 8.4% of their total wealth in PE at the time of capital call. The rest of the policies are hardly affected: consumption increases only slightly (Figure 3, Panel B), and the stock-bond split is unaffected (Panel D).

Along with the increase in allocation, there is a corresponding welfare gain of 1.25% of total wealth (Table 2, Panel A). This amount is large: it corresponds to 15% of the amount committed to PE. Equivalently, the investor is willing to give up a return premium of 1.10% out of PE's expected return, forever.

The intuition for a high commitment-quantity risk premium is that the investor's welfare declines when their portfolio moves away from the optimal value. There are two reasons for the welfare loss. The first is anticipation during the commitment period that the portfolio will continue to be sub-optimal in the holding period. At that time, the investor cannot consume out of illiquid wealth, and so a sub-optimal portfolio reduces consumption and portfolio volatility leads to consumption volatility (Section B.2). Second, fees during the commitment period are proportional to committed capital and taken out of liquid wealth; thus fees are volatile and a sub-optimal portfolio has a direct liquid wealth consequence.[23] Allowing the investor to choose their commitment at the time of a capital call enables them to avoid the impact of public market movements.

---

[23]We explore the relative contribution of fees to this effect with a sensitivity analysis in Section A.

34

C.3. Interaction between Timing and Quantity Risks

We next examine the interaction between timing and quantity risks by solving the investor's problem in Economy 3: deterministic commitment period with the ability to adjust commitments at the time of the capital call.

Commitment-quantity risk and commitment-timing risk are related but distinct. The intuition is as follows. In the presence of quantity risk, timing risk induces uncertainty on the distribution of the PE allocation and thus on the amount of quantity risk. However, we have shown that convex discounting pulls the welfare cost of timing risk down to nearly zero. As a result, the agent is indifferent between a deterministic and a random time of capital call. In the absence of quantity risk, the distribution of the PE allocation during the commitment period becomes irrelevant, but timing risk still induces uncertainty on the delay until when PE returns will be earned. Is the investor willing to pay to make this delay deterministic?

In Economy 2, quantity risk is turned off but timing risk is on. We compare this economy to Economy 3 where they are both turned off. Both allocations and costs are similar: The welfare difference between these two economies is almost zero, with 0.02% welfare cost and 0.01% return premium (Table 2, Panel A).

We conclude that timing risk carries a negligible premium irrespective of whether quantity risk is present in the economy. In all cases, the investor values the opportunity to adjust their PE allocation at the time of the call to avoid that it becomes too large.

C.4. Removing the Commitment Period

Next we remove the commitment delay (Economy 4), so that the investor's capital is invested as soon as it is committed. This allows us to make two useful comparisons. First, by comparing Economy 4 to the baseline, we can assess the overall effect of capital commitment on asset allocations and welfare. Second, by comparing Economies 2 and 4 (both with no commitment-quantity risk), we can isolate the impact of the commitment delay, i.e., the specific effect of not earning PE returns during the commitment

period.

Interestingly, the optimal portfolio allocation changes between Economies 0 and 2, but not between Economies 2 and 4. This means that the commitment delay has no effect on the allocation once quantity risk has been removed. The investor increases their consumption slightly, in anticipation of the distributions that they will now receive earlier. The stock-bond split does not change. Thus, it is commitment-quantity risk, not the commitment delay, that generates the under-allocation to PE.

However, removing the commitment delay has a large effect on welfare. The investor is willing to pay 2.8% of their wealth to switch from Economy 0 to Economy 4. This is more than twice the amount they are willing to pay to switch from Economy 0 to Economy 2, i.e. to adjust their allocation at call time. Equivalently, the return premium associated with moving from Economy 0 to Economy 4 is 2.4%, also more than twice that of moving to Economy 2. These effects are due to the increase in active investment time: more time is spent with assets invested in PE, as opposed to waiting in the commitment period.

C.5. Distribution Timing Risk

We now draw a comparison between timing risk applied to commitment and applied to distribution. To do so, we solve Economy 5, which has commitment risk, but a deterministic holding period: the time between the capital call and distribution is fixed at $\frac{1}{\lambda_D}$.

We find that distribution-timing risk is meaningful, and substantially more important than commitment-timing risk, but less impactful than commitment-quantity risk. Table 2 shows that the optimal PE commitment increases from Economy 0 to Economy 5, from 5.2% to 6.8%. The investor also increases their consumption during the commitment period (Figure 3, Panel A), and the PE allocation at which the investor strategically defaults is higher than in the baseline economy. However, the stock-bond allocation remains similar to the one in the baseline economy.

Distribution timing risk is costly: The investor is willing to pay an initial

36

welfare cost of 0.4% to remove this risk, or equivalently to accept a permanent decrease in the PE fund return of 0.4% as well. This result is in line with what is documented in Ang, Papanikolaou, and Westerfield (2014): The investor prefers certainty in the timing of distributions. However, both the welfare cost and the return premium are less than half those of commitment-quantity risk.[24]

The cost of distribution-timing risk is different from that of commitment-timing risk. During the commitment period, the investor is only exposed to public market volatility whereas during the holding period, they are exposed to both public and private market volatility. As a result, it is more volatile (Figure 2, Panel B) and the investor is more sensitive to the duration of the holding period than to the duration of the commitment period. Certainty on the distribution timing allows them to better handle their consumption stream and is thus preferred.

V.   Sensitivity Analysis

In this section, we modify the model parameters both to assess the robustness of the results we just presented – the under-allocation to PE, the large welfare cost and return premium of commitment-quantity risk – and to better understand the mechanism behind these results.

A.   Default and Fees

The investor has the option to strategically default any time during the commitment period and at capital call. In the event of default, the investor is no longer obligated to provide the committed capital and they do not have to pay the fees associated with their commitment any more. The cost

---

[24]The discrepancies between our values of welfare costs and the ones presented in Ang, Papanikolaou, and Westerfield (2014) have two sources. First, the Ang, Papanikolaou, and Westerfield (2014) model does not have a commitment period. Second, our parameters are chosen from a detailed calibration to private equity data described in Section A.

of strategic default is that the investor is being banned from PE, so their investment opportunity set reduces to the stock and bond. In this section, we analyze the value of the investor's option to default.

In our baseline economy, the investor defaults as soon as their PE allocation reaches 20.8% of total wealth. This threshold is well below the allocation to the bond (40.5%) so the investor would have enough cash to pay for the capital call. Despite the early default, it is a near zero-probability event (Section B). Thus, commitment default occurs early, but with very low probability.

We test whether this result is specific to our calibration by varying parameters such that the PE allocation can become large during the commitment period. We examine four cases: (i) the stock has low expected returns, (ii) the stock has high return volatility, (iii) the investor has a low risk aversion, and (iv) the commitment period is longer. In all scenarios, strategic default occurs early and remains a near zero-probability event as shown in Table 4.

To understand why the investor does not wait for the capital call to default, we vary the level of fees that the investor must pay to maintain their commitment. Figure 4, Panel A shows the investor's value function during the commitment period as a function of the PE allocation and the fees. The convexity is created by the option to default, and the point of default is at a lower allocation (ξ) when the level of fees is higher. As we expect, the investor defaults later when fees are lower; lower fees imply that both the immediate payments based on commitment are lower, and that the overall value of private equity is higher (the cost of losing future investment opportunities by defaulting is higher).

We next compare the point at which the investor would default while waiting for a capital call to the point of default at the moment of capital call (Figure 4, Panel B). When fees are near zero, the investor always waits for the call to default, but, when fees are high, the investor defaults early.

The fees paid during the commitment period are proportional to the PE commitment. Thus, public market declines have a similar impact on the PE allocation and on the fees paid: the allocation and the fees, relative to liquid wealth, increase. To test whether the point of strategic default is driven by

38

high fees in extreme states of the PE allocation, we set an upper threshold on the fees paid as a fraction of liquid wealth. Thus, fees are proportional to commitment capital up to this threshold. Figure 4, Panel C shows that the default point does not change much when decreasing the cap from 0.40% to 0.15%[25]: the default point increases from 20.8% (the default point in our baseline economy) to 26.2%. More generally, the value function does not change much, and the optimal PE commitment remains the same. Hence, for high PE allocations, welfare is not driven by fees but by the investor's portfolio allocation.

We thus have contrasting results. On the one hand, eliminating fees in states with a large allocation to PE does not affect welfare much. In these states, the investor's value function is mostly driven by the portfolio allocation. On the other hand, fees have a substantial impact on welfare (Figure 4, Panel A) for PE allocations that are close to the optimal commitment. As a result, they have a substantial effect on the investor's default point and on their optimal PE commitment. Thus, default is heavily influenced by fees because reducing fees raises the value of future PE investments – what is given up in default.

B. Secondary Market

The secondary market offers the possibility to exit an allocation that has become too large. However, in the baseline economy, the investor only infrequently makes secondary market sales. This result is not sensitive to the secondary market haircut. Eliminating the haircut on the secondary market hardly changes this result. The investor's commitment increases only slightly, from 5.2% to 5.7% (Table 3, Panel A), which remains far from the optimal level upon capital call. Similarly, increasing the haircut to 40% only decreases the allocation to 5% (Panel B).

---

[25]A cap of 0.15% (resp., 0.2% and 0.4%) means that the investor pays the minimum of 2% of commitment capital or 0.15% of liquid wealth. This implies that fees are not paid on commitments above 7.5% (resp., 10% and 20%).

The reason for this infrequent use is that, as in practice (see Section E), the secondary market forces the investor to sell their entire stake in a fund and wait until the next fund is raised and then until capital is deployed, thereby losing PE excess returns.

The secondary market has a larger *indirect* effect: it raises the premiums associated with capital commitment, particularly the premiums associated with quantity risk and the commitment delay (Table 3, Panels A and B). The welfare cost of commitment-quantity risk increases from 1.2% to 1.4% as the haircut declines from 13.8% to 0%. The willingness to pay to eliminate the commitment period increases as well, with a welfare cost going from 2.8% to 3.25%. The return premiums similarly increase. Furthermore, improving the liquidity of the secondary market increases welfare more when other liquidity frictions are removed (Table 3, Panel C). For example, the investor is willing to give up 0.07% of their wealth for a liquid secondary market in the baseline economy, but 0.25% when they can freely adjust their PE allocation at a capital call.

Putting the two results together, we conclude that the two types of liquidity – commitment risk and ease to sell on the secondary market – are complements, not substitutes. Increasing liquidity along the first dimension increases the willingness to pay to remove frictions along the second dimension. Moreover, the indirect effect of the secondary market on the welfare cost of capital commitment is larger than the direct welfare effect of the secondary market itself.

Thus, our model implies that the development of a PE secondary market increases the investor's desire to alleviate commitment-quantity risk, rather than satiate that desire. Intuitively, the investor is willing to invest more in PE if they have an easier way out, so this out becomes more valuable when commitment-quantity risk is absent.[26]

_____

[26]All these results are also robust to allowing the investor to reinvest in a new fund as soon as they sell their position on the secondary market, instead of waiting for the fund's end of life (non-tabulated).

C. Calibrated Parameters

We examine the sensitivity of the optimal PE commitment to the model parameters, and benchmark the variations to the ones that would be observed in a fully liquid model. Results are shown in Figure 5.

Quite strikingly, for all the parameters except the correlation, PE allocations are less sensitive to parameter changes in our model than they are in a fully liquid model. When PE expected return is below 12%, the optimal commitment to PE is zero (Panel A). Above that threshold, the PE commitment increases linearly with expected returns: each additional percentage of expected return increases the allocation by 2.25%. However, if PE were a fully liquid asset, the increase would be nearly twice as fast: 4.2% for each additional 1% return.

Similarly for PE volatility, in the baseline model, varying PE volatility from 35% to 25% changes the optimal commitment from 4% to 17.9% (Panel B). In contrast, in the fully liquid model, the corresponding allocation changes from 11.7% to 38.3%. We obtain similar results when varying stock's expected return (Panel C), stock volatility (Panel D), and the investor's risk aversion level (Panel F). The sensitivity of the allocation to changes in the correlation is low, but it was also the case in the Merton economy (Panel E).[27]

These differences in the sensitivity of the optimal commitment in our model compared to the fully liquid case are consistent with the hypothesis that it is stochastic illiquidity that creates an issue for the investor. Changing parameters so as to make PE more attractive pulls up the optimal allocation. However, because the investor cannot freely enter or exit their investment, they are less willing to take larger positions because there is an increased chance that their portfolio moves significantly away from the optimum, making their consumption more volatile.

---

[27]The correlation result is consistent with Ang, Papanikolaou, and Westerfield (2014), who examine a model of illiquidity without commitment, and the risk aversion result is similar to the result in Giommetti and Sorensen (2020).

To illustrate, consider the case of an increase in PE expected returns from $\nu = 14\%$ to $\nu = 18\%$, the optimal commitment goes from 5.2% of wealth to 14.4%, and the probability that consumption decreases by 5% (relative to liquid wealth) during the holding period, assuming it is initially at its optimum, increases from 1.9% to 11.3%. In contrast, in a fully liquid Merton economy, the optimal allocation increases from 14.2% to 30%, but consumption, relative to liquid wealth, is maintained at its optimum. Illiquidity means that smaller changes in PE allocation are associated with more consumption volatility because the investor can only consume out of liquid wealth.

Changing the intensities of capital calls and distributions has moderate effects on the investor's optimal commitment. Increasing λ_C makes the commitment period shorter on average and less variable. This effect reduces both commitment risk and commitment delay and results in a larger PE commitment. Panel G shows that the optimal PE commitment more than doubles when the average commitment period changes from five years ($\lambda_C = 0.2$) to one year ($\lambda_C = 1$): 3.1% to 7.6%. In the limit of λ_C going to infinity, the optimal PE commitment goes to 8.3%, which is the allocation in Economy 4, when the investment is made at fund inception (no commitment period).

The sensitivity of the optimal PE commitment to the intensity of capital distributions is hump-shaped (Panel H): it is slightly increasing for holding periods that are longer than 6.5 years (i.e., $\lambda_D = 0.15$), and decreasing for holding periods that are shorter than 6.5 years. This hump shape is the result of two opposite effects. With a low λ_D – a long holding period – the PE allocation is highly variable over time; hence the investor pulls down their commitment. As λ_D increases, this variation is lower (the holding period is shorter), hence PE commitment increases. With a high λ_D, the active investment time is short, so the investor is not willing to pledge a large amount because that amount will not be optimal in the holding period.

D. PE Under-Allocation, Welfare Costs and Return Premiums

Our first key findings: the significant under-allocation to PE due to commitment-quantity risk and the premium associated with commitment-quantity risk, are robust and stable with respect to changes in underlying return and preference parameters (i.e., parameters other than λ_C and λ_D). Figure 5, Panels A to F, show the under-allocation – the difference between the optimal allocation in Economy 2 (choose commitment on call, dashed line) and the optimal commitment in the baseline economy (solid line) – as we vary these return and preference parameters. Similarly, Figure 6, Panels A to F, display the sensitivity of the return premium of commitment-quantity risk (dashed lines). We observe little change when varying these parameters.

This finding does not hold for the welfare cost of commitment-quantity risk (solid lines on Figure 6, Panels A to F): The investor is willing to pay more to eliminate quantity risk when PE expected returns increase, or when either PE volatility, correlation with the stock returns, or risk aversion decrease. The welfare cost is indeed a one-time payment at time 0 for the total risk that the investor will bear, at $t = 0$. As such, it is mechanically increasing in the optimal PE commitment as described in Section F; see Figure 5. In contrast, the return premium represents a cost *per unit of PE allocation* and is much less sensitive to the optimal PE commitment.

The two intensity parameters, (λ_C, λ_D), have large effects, both on the magnitude of the PE under-allocation and on the return premium of commitment-quantity risk. These two parameters drive the expected length of the commitment period relative to the overall fund life: a larger λ_C or a smaller λ_D yields a shorter average commitment period relative to fund life. In this case, there is less commitment risk, hence both the PE under-allocation and the return premium associated with commitment-quantity risk decrease (Figure 6, Panels G and H). Changes in the welfare cost are smaller because they are dampened by increases in the optimal PE commitment.

In our calibration, we have used the NAVs reported at the end of year 2020, for funds raised between 1991 and 2015. To have a sense of the im-

43

portance of NAV biases[28] on our calibration, we repeated the calibration with funds raised up until 2012. The sample size decreases significantly, but NAVs have a negligible impact on fund performance in that sub-sample. We find that the implied correlation between public and private equity decreases slightly in this sub-sample, from 66% to 63.5%. One possible explanation is that major macroeconomic shocks faced by funds raised in 2013-2015 were particularly similar to those faced by public equity (e.g. Quantitative Easing). This small change in the correlation does not affect our results.

All our results still hold with Epstein-Zin preferences, as shown in Internet Appendix IA.C. Varying the intertemporal elasticity of substitution affects the investor's optimal consumption, but the changes to the allocations in the risky assets are negligible, and so are the welfare costs and return premiums of commitment risk.

---

[28]NAVs result from subjective judgments about the appropriate valuation technique and input parameters for each portfolio company. Despite accounting rules and although NAVs have no direct impact on investors' wealth, fund managers may purposely smooth NAVs with the aim of facilitating investor relationship management (e.g. avoid negative return news), facilitating fund-raising (Barber and Yasuda (2017); Brown, Gredil, and Kaplan (2019)), or because they believe that public market returns are excessively volatile. Crain and Law (2016) provide evidence that NAVs are quite accurate overall, although sluggish. Nadauld, Sensoy, Vorkink, and Weisbach (2019) show that some secondary market transactions are executed at prices that differ significantly from NAV.

VI. Commitment Quantity Risk & Liquidity Cycles

In this section, we examine how time varying liquidity impacts our results. The key idea is that capital calls and distributions are correlated with returns. In principle, cycles can exacerbate commitment risk if PE funds demand liquidity when it is most difficult to provide it. However, cycles may also moderate commitment risk if PE expected returns are higher when capital is called more quickly.[29]

A. Model and Calibration

In our baseline model, the Poisson processes triggering capital calls and distributions are independent of each other and independent of PE fund returns. However, our setup enables us to relate the intensity of capital calls and distributions to expected returns.

In this section, we assume that the economy can be in one of two states $s_t = \{L, H\}$. State L corresponds to periods of low liquidity and state H to high liquidity. The state of liquidity s_t follows a continuous time Markov process with a transition probability matrix between t and $t + dt$ given by

$$M = \begin{pmatrix} 1 - \chi^L \, dt & \chi^L \, dt \\ \chi^H \, dt & 1 - \chi^H \, dt \end{pmatrix}. \tag{21}$$

We identify years of low liquidity as the vintage years that experienced the lowest intensity of capital calls. Specifically, we calculate the fraction of

---

[29]Part of the motivation for such an analysis in based on anecdotal evidence such as what Leibowitz and Bova (2009) report in 2008: "The horrendous declines presented liquidity problems even for many portfolio managers who were long-term oriented, had modest payment schedules, and a seemingly ample percentage of liquid assets. This perfect liquidity storm, layered on top of a perfect asset storm, resulted from a toxic combination of: 1) a need to fulfill prior commitments to private equity, venture capital, real estate, and hedge funds, 2) reduced distributions from these asset classes..."

capital that is called each quarter for each fund. We remove an age fixed effect, as in Robinson and Sensoy (2016), and average the residual quarterly intensity across all funds with a given vintage year. We rank the vintage years by their associated average intensity of capital calls. The lowest 30% are the vintage years 1992, 1993, 1999, 2000, 2002, 2007, 2008 and 2009. In line with intuition, these years include the burst of the dot-com bubble and the financial crisis.

We calibrate the stock return and volatility in the low liquidity state as the average S&P 500 log return and volatility over the five years starting with each low-liquidity year. We find $\mu^{low} = 6.63\%$ and $\sigma^{low} = 18.59\%$.

The other parameters are calibrated as described in Section A, but keeping only PE funds that have a low-liquidity vintage year. We obtain $\nu^{low} = 11\%$, $\phi^{low} = 37.5\%$, $\lambda_C^{low} = 0.34$ (commitment period of 2.9 years) and $\lambda_D^{low} = 0.14$ (holding period of 7.1 years).

We set the parameters in the high liquidity state so that their state-weighted average match their value absent cycles. Finally, in line with Nadauld, Sensoy, Vorkink, and Weisbach (2019), we set the secondary market haircut to 28% in the low liquidity state and 9% in the high liquidity state. Parameter values are summarized in Table 5.

The low liquidity state is therefore characterized by lower expected returns, higher volatility and a higher correlation between public and private equity. In addition, intensities of capital call and distribution are lower than in the high liquidity state, but the intensity of capital call is less sensitive than the intensity of capital distributions. This result is in line with the key insight of Robinson and Sensoy (2016): Net cash flows are procyclical because distributions are more procyclical than calls.

B.   Time-Varying Liquidity & Commitment Quantity Risk

We compare the value function before and after a call in the no-cycle economy (Figure 7, Panel A) to the value functions in the economy with cycles, in both liquidity states (Panel B). Points A and B are the optimal PE commitments in the low and the high state. Points C and D are the optimal

PE allocations during the holding period. A clear difference is that the presence of cycles induce a level shift (upwards) of the value function in both liquidity states, even though the probability-weighted average parameters are the same in the economy with and without cycles. This level shift is driven by the changes in the moments of the public equity returns, and not by the changes in the moments of the private equity returns. Indeed, Panel C shows that using the calibrated cycles parameters for the stock return moments, but setting all other parameters to their no-cycle values, suffices to produce the level shift. If, instead of having cycles on the stock return moments, we have them on PE moments, we do not observe any shift (Panel D).

Naturally, the level shift only affects the overall value function; it does not affect the optimal allocation and commitment. Cycles, however, also change the slope of the value function, and hence the optimal allocation and commitment. In addition, we observe that the change in the slope is similar in both liquidity states, implying that the optimal PE commitment and allocation in the high and low state are relatively close to each other. As shown in Table 6, the investor commits less to PE (4.3% in the low liquidity state and 4.7% in the high liquidity state) than in the economy without cycles (5.2%). As a result of the lower PE commitments, the welfare costs of commitment-quantity risk is also lower than in the no-cycle economy (see Section F). However, the return premiums of commitment-quantity risk are similar to the economy without cycles, in both liquidity states.

We test whether this result still holds when changing the impact of cycles on returns. Our cycle calibration was performed with realized returns, so returns are lower in the low liquidity state. But Haddad, Loualiche, and Plosser (2017) suggest that expected returns to both public and private markets are higher in the low liquidity state. We swap both public and private equity return parameters from the high state to the low state (Panel E). In the low (high) liquidity state, volatilities are high (low) but now, so are expected returns. These returns therefore act as an insurance: low liquidity and high volatility are paired with high expected returns. As a result, the value function does not change much between the low and the high liquidity states. Importantly, our previous results still hold: the slope shift is

47

here again similar in both states, and the return premiums of commitment-quantity risk are similar, in both states, to the one without cycles (Table 6, line 'both switched').

We next study variations of the calibrated parameter set, in which private and public equity returns are not both higher in the same state. We fix public equity expected returns at $\mu = 0.08$ in both high and low states; all the other parameters are set to their calibrated values. With this change, the stock return is higher in the low liquidity state than in our calibration. This is the state in which the investor has high needs in liquid wealth as there are fewer calls but much fewer distributions. There is therefore a new mismatch between stock returns and the investor's need in liquid wealth in the low liquidity state. The resulting slope shift leads to a much lower optimal PE commitment in the low state (Panel F), and a slightly higher PE commitment in the high state (Table 6, line 'μ fixed'). The large difference between these optimal commitments means that the investor would like to change their PE allocation each time the state changes. This additional risk slightly increases the return premium in *both* states.

Swapping the stock return from the high liquidity state to the low liquidity state exacerbates this result (Table 6, line 'μ switched'). Swapping the PE return from the high liquidity state to the low liquidity state but not the stock confirms it.

To conclude, we find that adding liquidity cycles has a large effect on welfare, a smaller effect on portfolio allocation, and an even smaller effect on commitment risk premiums. The impact of time-varying liquidity on commitment risk premiums depends on the *relative* effect on private and public markets.

VII. Is Commitment Risk Diversifiable?

In this section, we study whether commitment-quantity risk is diversifiable. On the one hand, by spreading their allocation across several funds, the investor may smooth their capital inflows and outflows.[30] It also allows the investor to sell only part of their PE exposure on the secondary markets. On the other hand, the investor risks a funding mismatch if earlier investments pay out late while later investments call early. In addition, much of the commitment risk premiums are driven by public market movements, which affect the allocations to all PE funds through the denominator effect.

A. Two Funds

We solve the two-fund model described in Section E and show results in Figure 8 and Table 7. The optimal commitment allocates 4.2% to each fund (8.5% total), as opposed to a commitment of 5.2% when a single fund is available.

The first result is that access to a second fund is not particularly valuable. Diversification increases investor's welfare, but modestly: the investor is willing to give up 1.0% of their wealth, or accept a permanent reduction in PE returns of 0.9% in order to gain access to a second fund (Table 7, Panel C). For comparison, the investor would be willing to give up 1.25% of their wealth, or a return premium of 1.1% to alleviate commitment-quantity risk in the one-fund model. In other words, despite the gains from cash

---

[30]This assumption is supported empirically by Robinson and Sensoy (2016): "Most variation in fund-level cash flows is purely idiosyncratic across funds of a given age at a given point in time (...) this suggests that liquidity shocks arising from the uncertain timing of calls and distributions can be significantly mitigated by holding a portfolio of investments diversified both across different funds of the same age and across funds of different ages. For example, for buyout funds the standard deviation of quarterly net cash flows averages 11.57% of committed capital, and this standard deviation shrinks to 4.54% in a portfolio of all buyout funds in the sample."

flow diversification, the investor prefers to control their quantity invested by adjusting their commitment rather than through use of a second fund.

Consistent with this result, the investor does not do much liquidity management – their allocation to a second fund does not depend much on their allocation to the first fund. If the investor has a large allocation, either a commitment or a holding, in fund 1, they will commit less to fund 2 (see Figure 8, Panel A). However, this effect is quite small: If the allocation to fund 1 is 8.8%, twice the optimal commitment to the first fund, the investor commits 3.8% to the second fund instead of 4.4%.

In contrast with the one-fund case, a capital call does not always trigger a utility gain with multiple PE funds. Figure 8, Panel D, shows that when the commitment to fund 1 is called, the welfare change is usually positive if the second fund is in the commitment period (dashed line) but negative if it is in the holding period (solid line). Indeed, if the second fund has not been called yet, there is no capital deployed in PE and liquid wealth is high (because of the previous fund's distributions). The capital call is hence *good news*. In contrast, if the second fund has already been called, the capital call of fund 1 occurs at a time of low liquid wealth. There is a *funding mismatch*, i.e., the capital call is *bad news*.

To evaluate whether diversification decreases commitment-quantity risk, we solve Economies 0, 2, and 4 with two PE funds; results are shown in Table 7 (Panel A). The optimal PE commitment increases compared to the one-fund case, but the investor still commits much less in Economy 0 than in Economy 2 (8.5% versus 11.4%). Hence, the under-allocation to PE remains strong.

Consistent with the increase in the investor's allocation, the welfare cost of commitment-quantity risk increases to 1.4% (versus 1.25% with one fund), but the return premium slightly decreases to 0.8% (versus 1.1% with one fund). We conclude that the investor only modestly benefits from liquidity diversification, even though we overstate its value by assuming purely idiosyncratic liquidity shocks.

To understand why the gains from diversification are small, we examine how the value function reacts to increases in PE allocation. In the economy

with one fund, we have shown that the key issue for the investor is the possibility of a decline in the stock price, causing the PE allocation to be too large (denominator risk). This problem is equally present with multiple funds. In Figure 8, Panel B, we keep the allocation to fund 2 at its optimal level and vary the allocation to fund 1 (thin line). The penalty for having an allocation to fund 1 that becomes too large (moving right) is relatively small and comparable to the penalty in the one-fund model. However, if the public market value declines, the allocations to both funds increase together. To reflect this joint effect, we force ξ^2 to vary together with ξ_1 ($\xi^2 = \xi^1$). The value function becomes much more concave, with a larger penalty when the allocation to fund 1 increases. These results are similar when both funds are in their holding period (Panel C). In other words, *denominator risk* cannot be diversified.

B.   *Infinite Number of Funds*

We now solve the infinite-fund problem described in Section E. In this setup, each fund is infinitely small and the two key variables are the aggregate PE commitment and the aggregate PE investment as a function of liquid wealth: $\pi = \frac{X^\infty}{W+Y^\infty}$ and $\xi = \frac{Y^\infty}{W+Y^\infty}$. Figure 9 plots the former as a function of the latter. The optimal aggregate commitment decreases rapidly as the aggregate investment increases: it reaches 0 for $\xi \geq 21\%$. At the steady state, the fraction of liquid wealth invested in PE is 19.6% and only 2.2% is committed (thus uncalled). Therefore, the total PE allocation amounts to 21.9%, which is a large increase compared to the one- and two-fund cases.[31]

In addition, the investor sustains a level of investment that is much higher than the level of commitment. This is possible with an infinite number of funds because the timing of capital calls and distributions is no longer stochastic. Thus, it is no longer possible to have an unanticipated funding mismatch. As a result, the investor is willing to give up 5.5% of their wealth,

---

[31]The increase in PE allocation as a function of the number of funds available is consistent with the empirical evidence given in Section C.

or accept a reduction of 3.0% in PE returns to have access to an infinite number of funds instead of two funds (Table 7, Panel C).[32]

However, we observe that the investor is still under-allocated to PE relative to what they would allocate without commitment risk (Table 7, Panel B). When given the possibility to choose-on-call, the PE allocation is 28.5%. As the change in allocation is large, the welfare premium associated with eliminating quantity risk is also large at 4.3% (versus 1.4% with two funds) and the return premium is only slightly smaller (0.7%) than with two funds (0.8%); see Panel C. The effects of diversification are therefore small because the denominator risk remains: undiversifiable public market movements still drive volatility in the PE allocation and thus excess volatility in consumption.

Surprisingly, the investor rarely uses the secondary market with an infinitely of funds, even though they can now sell any fraction of their aggregate investment. The reason is that they stop committing to new funds when the aggregate investment exceeds 21%. The invested amount then naturally decreases as capital is distributed (at rate $\lambda_D dt$). The larger λ_D, the less the investor needs the secondary market. Our model thus offers a unique insight regarding the secondary market: investors with few PE positions should be reluctant to use it because they need to liquidate a large chunk of their PE position if they do; and investors with many PE positions should be reluctant to use it because if they stop committing, their PE exposure decreases rapidly.

VIII. Conclusion

This paper proposes an optimal dynamic portfolio allocation model that includes capital calls and distributions with uncertain timing. We calibrate this model to data on PE fund cash flows and show that ex-ante commitment

---

[32]Investing in many funds requires the LP to manage relationships with many GPs, which is costly. If the running cost of having to manage these relationships is lower than a 3% reduction in PE returns, the total welfare gain of having access to many funds (compared to two funds) remains positive. We thank the referee for pointing this out.

has large effects on investors' portfolios and welfare.

One key finding is that investors want to change their capital allocations when capital is called – most often to increase the allocation. Hence, investors are under-allocated to PE, and willing to pay a large premium to adjust the quantity committed upon capital call. A direct implication of PE under-allocation is that the demand for changing the amount committed is asymmetric. Investors would rarely want to reduce their allocation to PE, but highly value the option to top-up their allocation when capital is called. For example, they would highly value the possibility to co-invest, which is consistent with the institutional details reviewed in Section E. We could hence see recent key institutional developments as allowing some investors to alleviate commitment-quantity risk.

With one fund, the return premium of commitment-quantity risk amounts to 1.25% of the investors' initial wealth, or equivalently to a permanent loss of 1.1% of private equity returns. It is larger than the premium investors are willing to pay to eliminate other liquidity frictions: timing uncertainty and the limited tradability of PE investments. Furthermore, commitment risk premiums increase with secondary market liquidity, and increasing the number of funds does not allow the investor to diversify commitment risk, particularly when liquidity is time varying.

A natural question that arises is "Why do LPs not simply diversify the problem away?" Making multiple different capital commitments reduces the size of each commitment and call, removing the lumps. We find that this diversification with multiple funds is only weakly helpful. Investors care about the fraction of their wealth that they have committed, and public market movements change all the allocations' denominators at once, so the key risk – commitment-quantity risk – cannot be diversified. Worse, without perfect and full diversification, there is always a potential funding mismatch in which holdings in one fund impact the welfare value, through future capital calls, of commitments in other funds.

REFERENCES

Ang, Andrew, Andrés Ayala, and William N. Goetzmann, 2018a, Investment beliefs of endowments, *European Financial Management* 24, 3–33.

Ang, Andrew, Bingxu Chen, William N. Goetzmann, and Ludovic Phalippou, 2018b, Estimating private equity returns from limited partner cash flows, *Journal of Finance* 73, 1751–1783.

Ang, Andrew, Dimitri Papanikolaou, and Mark M Westerfield, 2014, Portfolio choice with illiquid assets, *Management Science* 60, 2381–2617.

Arcot, Sridhar, Zsuzsanna Fluck, José-Miguel Gaspar, and Ulrich Hege, 2015, Fund managers under pressure: Rationale and determinants of secondary buyouts, *Journal of Finance* 115, 102–135.

Axelson, Ulf, Tim Jenkinson, Per Strömberg, and Michael S. Weisbach, 2013, Borrow cheap, buy high? the determinants of leverage and pricing in buyouts, *Journal of Finance* 68, 2223–2267.

Axelson, Ulf, Per Strömberg, and Michael S. Weisbach, 2009, Why are buyouts levered? the financial structure of private equity funds, *Journal of Finance* 64, 1549–1582.

Banal-Estañol, Albert, Filippo Ippolito, and Sergio Vicente, 2017, Default penalties in private equity partnerships, Working Paper.

Barber, Brad M., and Ayako Yasuda, 2017, Interim fund performance and fundraising in private equity, *Journal of Financial Economics* 124, 172–194.

Barrot, Jean-Noël, 2017, Investor horizon and the life cycle of innovative firms: Evidence from venture capital, *Management Science* 63, 2773–3145.

Bollen, Nicolas P. B., and Berk A. Sensoy, 2021, How much for a haircut? Illiquidity, secondary markets, and the value of private equity, *Financial Management* Forthcoming.

Brown, Gregory W., Oleg Gredil, and Steven N. Kaplan, 2019, Do private equity funds manipulate reported returns?, *Journal of Financial Economics* 132, 267-297.

Brown, Gregory W., Robert S. Harris, Wendy Hu, Tim Jenkinson, Steven N. Kaplan, and David T. Robinson, 2021, Can Investors Time Their Exposure to Private Equity?, *Journal of Financial Economics* 139, 561–577.

Cavagnaro, Daniel R., Berk A. Sensoy, Yingdi Wang, and Michael S. Weisbach, 2019, Measuring institutional investors' skill at making private equity investments, *Journal of Finance* 74, 3089–3134.

Cox, John C., and Chifu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 49, 33 – 83.

Crain, Nicholas G., and Kelvin K. F. Law, 2016, Spreading sunshine: Fair value accounting and information production, Working Paper.

Da Rin, Marco, and Ludovic Phalippou, 2017, The importance of size in private equity: Evidence from a survey of limited partners, *Journal of Financial Intermediation* 31, 64–76.

Dai, Min, Peifan Li, Hong Liu, and Yajun Wang, 2015, Portfolio Choice with Market Closure and Implications for Liquidity Premia, *Management Science* 62, 368–386.

Degeorge, François, Jens Martin, and Ludovic Phalippou, 2016, On secondary buyouts, *Journal of Financial Economics* 120, 124–145.

Dimmock, Stephen G., Neng Wang, and Jinqiang Yang, 2023, The Endowment Model and Modern Portfolio Theory, *Management Science* Forthcoming.

Dybvig, Philip H., and Chifu Huang, 1988, Nonnegative wealth, absence of arbitrage, and feasible consumption plans, *Review of Financial Studies* 1, 377–401.

Fang, Lily, Victoria Ivashina, and Josh Lerner, 2015, The disintermediation of financial markets: Direct investing in private equity, *Journal of Financial Economics* 116, 160–178.

Gârleanu, Nicolae, 2009, Portfolio choice and pricing in illiquid markets, *Journal of Economic Theory* 144, 532–564.

Giommetti, Nicola, and Morten Sorensen, 2020, Optimal Allocation to Private Equity, Working paper.

Haddad, Valentin, Erik Loualiche, and Matthew Plosser, 2017, Buyout Activity: The Impact of Aggregate Discount Rates, *Journal of Finance* 72, 371–414.

Kahl, Matthias, Jun Liu, and Francis Longstaff, 2003, Paper millionaires: how valuable is stock to a stockholder who is restricted from selling it?, *Journal of Financial Economics* 67, 385–410.

Korteweg, Arthur G., 2019, Risk Adjustment in Private Equity Returns, *Annual Review of Financial Economics* 1, 131–152.

Korteweg, Arthur G., and Mark M. Westerfield, 2022, Asset allocation with private equity, *Foundations and Trends in Finance* 13, 95–204.

Leibowitz, M., and A. Bova, 2009, Portfolio Liquidity, *Morgan Stanley Research Portfolio Strategy* .

Lerner, Josh, Jason Mao, Antoinette Schoar, and Nan R. Zhang, 2022, Investing Outside the Box: Evidence from Alternative Vehicles in Private Capital, *Journal of Financial Economics* 143, 359–380.

Lerner, Josh, and Antoinette Schoar, 2004, The illiquidity puzzle: theory and evidence from private equity, *Journal of Financial Economics* 72, 3–40.

Longstaff, Francis A., 2001, Optimal Portfolio Choice and the Valuation of Illiquid Securities, *Review of Financial Studies* 14, 407–431.

Longstaff, Francis A., 2009, Portfolio Claustrophobia: Asset Pricing in Markets with Illiquid Assets, *American Economic Review* 99, 1119–1144.

Metrick, Andrew, and Ayako Yasuda, 2010, The Economics of Private Equity Funds, *Review of Financial Studies* 23, 2303–2341.

Nadauld, Taylor D., Berk A. Sensoy, Keith Vorkink, and Michael S. Weisbach, 2019, The liquidity cost of private equity investments: Evidence from secondary market transactions, *Journal of Financial Economics* 132, 158–181.

Phalippou, Ludovic, 2021, *Private Equity Laid Bare* (Independent).

Robinson, David T, and Berk A Sensoy, 2016, Cyclicality, performance measurement, and cash flow liquidity in private equity, *Journal of Financial Economics* 122, 521–543.

Sorensen, Morten, Neng Wang, and Jinqiang Yang, 2014, Valuing Private Equity, *Review of Financial Studies* 27, 1977–2021.

Table 1. Calibrated Parameters

This table displays the values of parameters obtained from the calibration described in Section A.

Panel A: Calibrated parameters in the model with a finite number of funds

Parameter	Symbol	Parameter value
Risk-free rate	r	0.03
PE expected returns	ν	0.14
Stock expected returns	μ	0.08
PE fund volatility	ψ	0.335
Stock volatility	σ	0.150
Correlation stock & PE	ρ_L	0.66
Correlation between PE funds	ρ_{PE}	0.68
Intensity of capital call	λ_C	0.344
Intensity of capital distribution	λ_D	0.174
Secondary market haircut	α	13.8%
Investor's time discounting	δ	0.05
Investor's risk aversion	γ	4
Fee on commitment	f	2%

Panel B: Calibrated parameters in the model with an infinite number of funds

Parameter	Symbol	Parameter value
Correlation between stock & PE portfolios	ρ_L^∞	0.80
PE portfolio volatility	ψ^∞	0.276

Table 2. Optimal PE Allocation & Liquidity Frictions

The baseline economy (E0) is our central model, i.e. an economy in which investors need to commit ex-ante on the amount that will be called, wait for capital to be called and receive capital calls and distributions at random times. These frictions are removed one at a time in Economies 1 through 4, as described in Section D. Economy 5 contains commitment risk, but the distribution time is deterministic. The PE allocation refers to the optimal PE commitment made at time τ_0 in Economies 0, 1 and 5 and to the optimal investment made at time τ_C in Economies 2, 3 and 4. As defined in Section F, the welfare cost is the amount investors are willing to pay to switch from E0 to a given economy, i.e. the willingness to pay to remove a given friction. The return premium is the additional return that PE should deliver (ν) in Economy 0 for the economy under consideration to be equivalent to E0, i.e. the return premium associated with a given friction. Results are shown with a secondary market haircut of 13.8%.

		PE allocation	Welfare cost	Return premium
E0	Baseline (All frictions)	5.23%		
E1	Deterministic call time	5.18%	0.01%	0.01%
E2	Choose quantity on call	8.32%	1.25%	1.10%
E3	E1 ∩ E2	8.32%	1.27%	1.11%
E4	No commitment period	8.32%	2.79%	2.41%
E5	Deterministic payout time	6.79%	0.43%	0.38%

Table 3. Secondary Market

This table illustrates the effect of the secondary market haircut on optimal PE allocations and on the welfare costs of commitment risk. Panels A and B report in column 1 the optimal PE allocations in Economies 0 to 5, for secondary market haircuts of 0% and 40%. Columns 2 and 3 report the cost of commitment-quantity risk. The welfare cost is the amount investors are willing to pay to switch from E0 to E2, i.e. the willingness to pay to remove commitment quantity risk. The return premium is the additional return that PE should deliver (ν) in Economy 0 for E2 to be equivalent to E0, i.e. the return premium associated with commitment quantity risk. Panel C reports the marginal impact of changing the secondary market haircut on investor welfare, in a given economy. The welfare cost of the default haircut, i.e., the amount that the investor is willing to pay to change this haircut from 13.8% to 0% (resp., from 13.8% to 40%) is displayed in the second (resp., third) column.

Panel A: Secondary market haircut of 0% (liquid market)

		PE allocation	Welfare cost	Return premium
E0	Baseline (All frictions)	5.67%		
E1	Deterministic call time	5.61%	0.01%	0.01%
E2	Choose quantity on call	9.62%	1.43%	1.33%
E3	E1 ∩ E2	9.62%	1.44%	1.35%
E4	No commitment period	9.71%	3.25%	2.97%
E5	Deterministic payout time	6.79%	0.36%	0.34%

Panel B: Secondary market haircut of 40% (illiquid market)

		PE allocation	Welfare cost	Return premium
E0	Baseline (all frictions)	4.98%		
E1	Deterministic call time	4.93%	0.01%	0.01%
E2	Choose quantity on call	7.80%	1.17%	0.99%
E3	E1 ∩ E2	7.80%	1.19%	1.01%
E4	No commitment period	7.73%	2.61%	2.18%
E5	Deterministic payout time	6.79%	0.47%	0.40%

Panel C: Welfare cost of changing the haircut in each economy

		Calibration (h=13.8%)	Liquid (h=0%)	Illiquid (h=40%)
E0	Baseline	0%	0.07%	-0.04%
E1	Deterministic call time	0%	0.07%	-0.04%
E2	Choose quantity on call	0%	0.25%	-0.13%
E3	E1 ∩ E2	0%	0.25%	-0.13%
E4	No commitment period	0%	0.54%	-0.23%
E5	Deterministic payout time	0%	0.00%	0.00%

Table 4. Scenario Analysis

This table illustrates the marginal impact of changing one parameter on the investor's optimal PE allocation (column 1), their bond allocation (column 2), and their probability of strategic default (column 3) in Economy 0 (baseline economy). Column 4 reports the welfare cost of commitment-quantity risk (column 4), obtained by comparing Economy 2 (the investor can update their commitment upon capital call) to the baseline economy. All other parameters are as in our calibration (see Table 1).

	PE alloc.	Bond alloc.	P(default)	Welfare cost
Baseline (E0)	5.23%	45%	$\approx 0\%$	1.10%
Low stock expected returns (μ=6%)	12.00%	67%	$\approx 0\%$	2.78%
High stock volatility (σ=0.20)	9.63%	69%	$\approx 0\%$	2.27%
Low risk aversion (γ=2)	10.65%	0%	$\approx 0\%$	0.75%
Long commitment period (λ_C=0.2)	3.27%	44%	$\approx 0\%$	1.41%

Table 5. Calibrated Parameters with Cycles

This table displays the values of parameters obtained in our extension of the model with liquidity cycles, following the calibration described in Section A. The values are given in the low and high liquidity states.

Parameter	Symbol	Low liquidity	High liquidity
Probability to switch from state H to L	χ^H	-	0.143
Probability to switch from state L to H	χ^L	0.333	-
Stock expected returns	μ	0.066	0.086
Stock volatility	σ	0.186	0.135
PE expected returns	ν	0.110	0.153
PE fund volatility	ψ	0.380	0.316
Correlation between stocks & PE	ρ_L	0.685	0.649
Intensity of capital call	λ_C	0.340	0.346
Intensity of capital distribution	λ_D	0.136	0.190
Secondary market haircut	α	28%	9%

Table 6. Optimal PE Allocation & Liquidity Frictions with Cycles

This table reports the impact of changing the return parameters of the stock and PE on the cost of commitment quantity risk. The welfare cost is the amount investors are willing to pay to switch from E0 to E2, i.e. the willingness to pay to remove commitment quantity risk. The return premium is the additional return that PE should deliver (ν) in Economy 0 for E2 to be equivalent to E0, i.e. the return premium associated with commitment quantity risk.

	μ^L	μ^H	ν^L	ν^H	Opt. comm. Low	High	Welfare cost Low	High	Ret. premium Low	High
no cycles	8%		14%		5.23%		1.25%		1.10%	
calibrated	6.6%	8.6%	11.0%	15.3%	4.33%	4.74%	1.07%	1.11%	1.09%	1.11%
both switched	8.6%	7.7%	15.3%	13.4%	5.34%	5.13%	1.24%	1.24%	1.09%	1.10%
μ fixed	8.0%	8.0%	11.0%	15.3%	3.96%	5.13%	1.15%	1.24%	1.14%	1.16%
μ switched	8.6%	7.7%	11.0%	15.3%	3.92%	5.50%	1.23%	1.35%	1.18%	1.20%
ν switched	6.6%	8.6%	15.3%	13.4%	6.27%	4.83%	1.31%	1.21%	1.15%	1.14%

61

Table 7. Multiple Funds

This table reports the benefits of increasing the number of funds in the investor's opportunity set. Panel A reports the total optimal PE commitment in Economy 0 (baseline economy; first line), Economy 2 (the investor can update their commitment upon capital call; second line) and Economy 4 (no commitment period; third line), when the investor has access to two private equity funds. The last two columns display the welfare costs and return premiums of Economies 2 and 4 compared to the baseline economy. Panel B reports the optimal PE aggregate commitment and investment, in the steady state, in Economies 0 and 2, as well. The steady state is defined as the state in which the expected change in invested capital is zero: $E[dX_t^\infty] = 0$, where X_t^∞ follows the dynamics given in equation (12). The last two columns display the welfare cost and return premium of commitment-quantity risk. Panel C reports the welfare costs and return premiums of having access to one fund instead of two, and two funds instead of an infinity.

Panel A: Two-fund allocations and costs of commitment-quantity risk			
	PE allocation	Welfare cost	Return premium
E0 Baseline economy	8.48%		
E2 Choose quantity on call	11.44%	1.37%	0.79%
E4 No commitment period	11.34%	3.39%	3.82%

Panel B: Infinite-fund allocations and costs of commitment-quantity risk				
	PE comm.	PE inv.	Welfare cost	Return premium
E0 Baseline economy	2.24%	19.62%		
E2 Choose quantity on call		28.50%	4.32%	0.74%

Panel C: Benefits of diversification		
	Welfare cost	Return premium
E0 Increasing from 1 fund to 2	0.98%	0.86%
E0 Increasing from 2 funds to an infinity	5.48%	3.02%

Figure 1. Model Validation

This figure illustrates the output of our calibration procedure, described in Section A. In Panel A, the blue marks represent the empirical fraction of called capital after n quarters since capital commitment, for n between 1 and 40. The red curve represents the model-implied fraction of capital calls for the our calibrated λ_C of 0.344. Panel B displays the empirical cumulative distribution function (cdf) of PMEs in our data sample (blue marks), and the model-implied cdf (red line). We consider the first payment (fee) as capital commitment date.

Panel A: Fraction of called capital

Panel B: Distribution of PME

Figure 2. Optimal Allocation and Policies in the Baseline Economy

Panel A represents the value function of the investor during the commitment period (dashed line) and the holding period (plain line) after adjusting for the jump in liquid wealth. Default is represented as a circle, sale on the secondary market as a square. Panel B displays the distribution of the PE allocation during the commitment period (dashed line) and holding period (plain line). Panel C displays the optimal consumption of the investor given their PE allocation. Panel D displays the optimal stock allocation. The vertical dotted line represents the optimal PE commitment ξ^*.
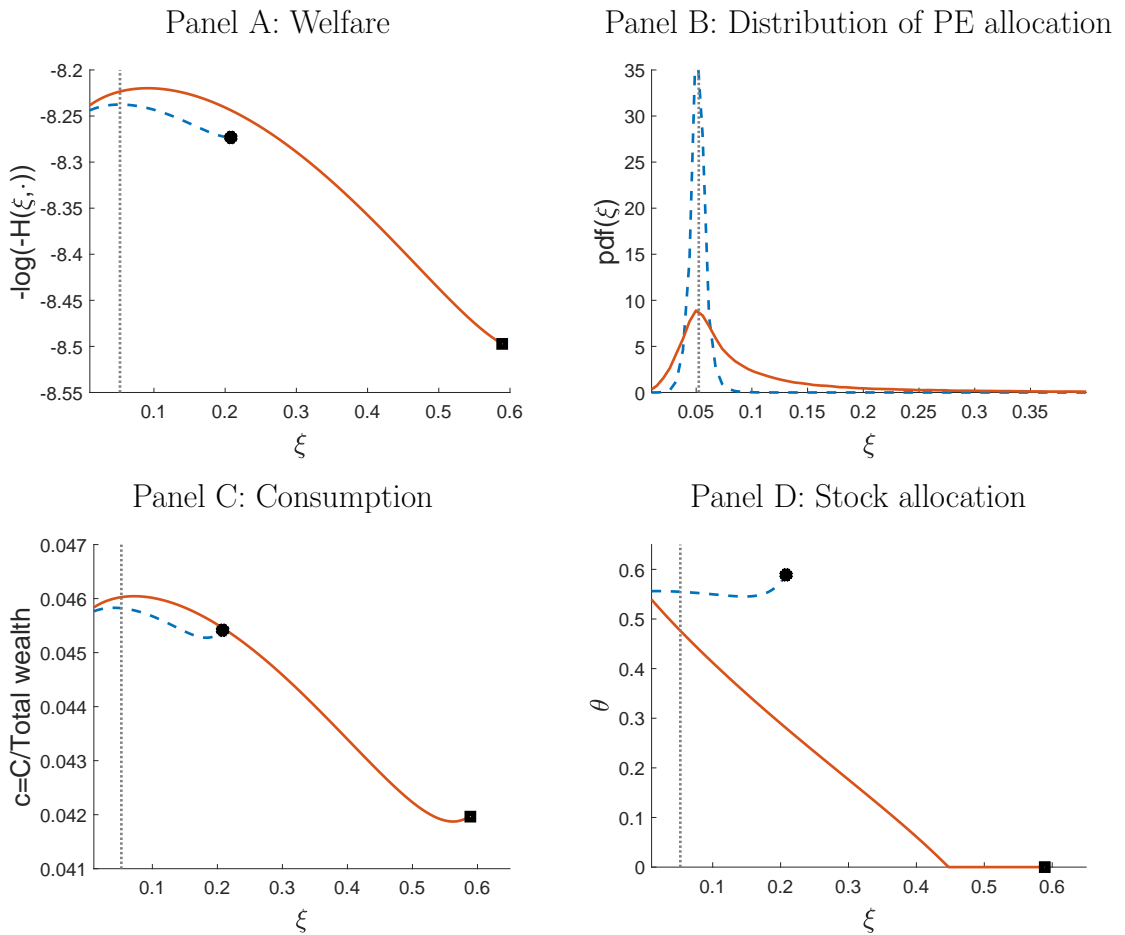
Figure 3. Consumption and Stock Allocation in Economies E0 to E5

This figure represents the optimal consumption rate and stock allocation before capital call (Panels A and C) and after capital call (Panel B), in the different economies. In Economies 2 to 4, the amount of capital invested in private equity is chosen at capital call therefore the relation between consumption (resp. stock allocation) and PE allocation before capital call is not displayed. After capital call, stock allocations overlap in economies E0 to E5.

Economies are summarized below:

Baseline economy E0	Model described in Section III	All risks on
Economy E1	Deterministic call time	Commitment-timing risk off
Economy E2	Choose quantity when called	Commitment-quantity risk off
Economy E3	Choose quantity when called + deterministic call time	Commitment risk off
Economy E4	No commitment delay	Commitment risk off
Economy E5	Deterministic payout delay	Distribution-timing risk off

Panel A: Consumption before call   Panel B: Consumption after call

Panel C: Stock allocation before call   Panel D: Stock allocation after call

65

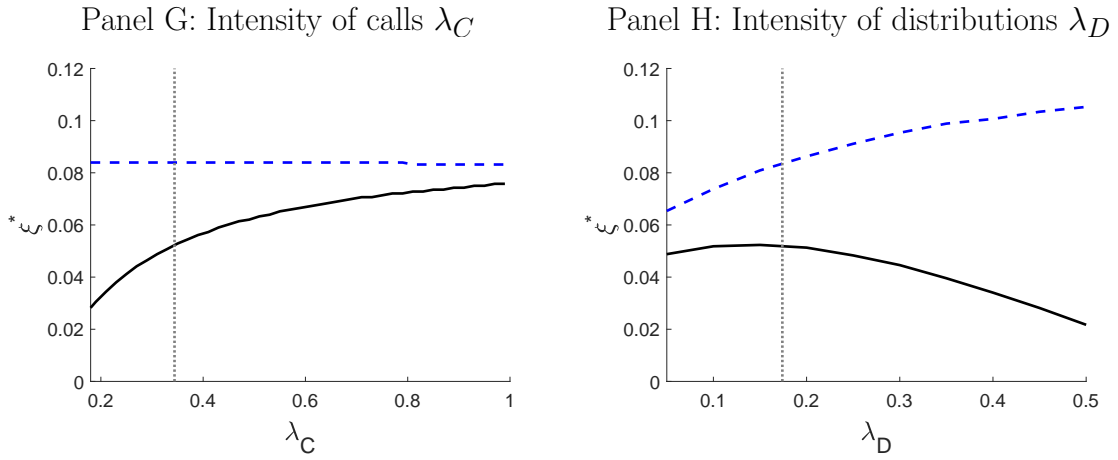**Figure 4.** Strategic Default in the Baseline Economy

Panel A displays the value function of the investor during the commitment period, for three levels of fees: 1%, 2% and 3% of the committed amount. The investor strategically defaults when their value function becomes lower than the value function in the Merton problem (no access to PE). Panel B displays the threshold allocation at which the investor strategically default during the commitment period, and at capital call, as a function of the fee level. Panel C displays the value function of the investor during the commitment period when fees are capped at three different thresholds: 0.15%, 0.2% and 0.4% of liquid wealth. When the fee (2% of the committed capital to PE) exceeds the threshold, it is set equal to the threshold.

Panel A: Value function

Panel B: Strategic default allocation

Panel C: Value function when capping fees

66
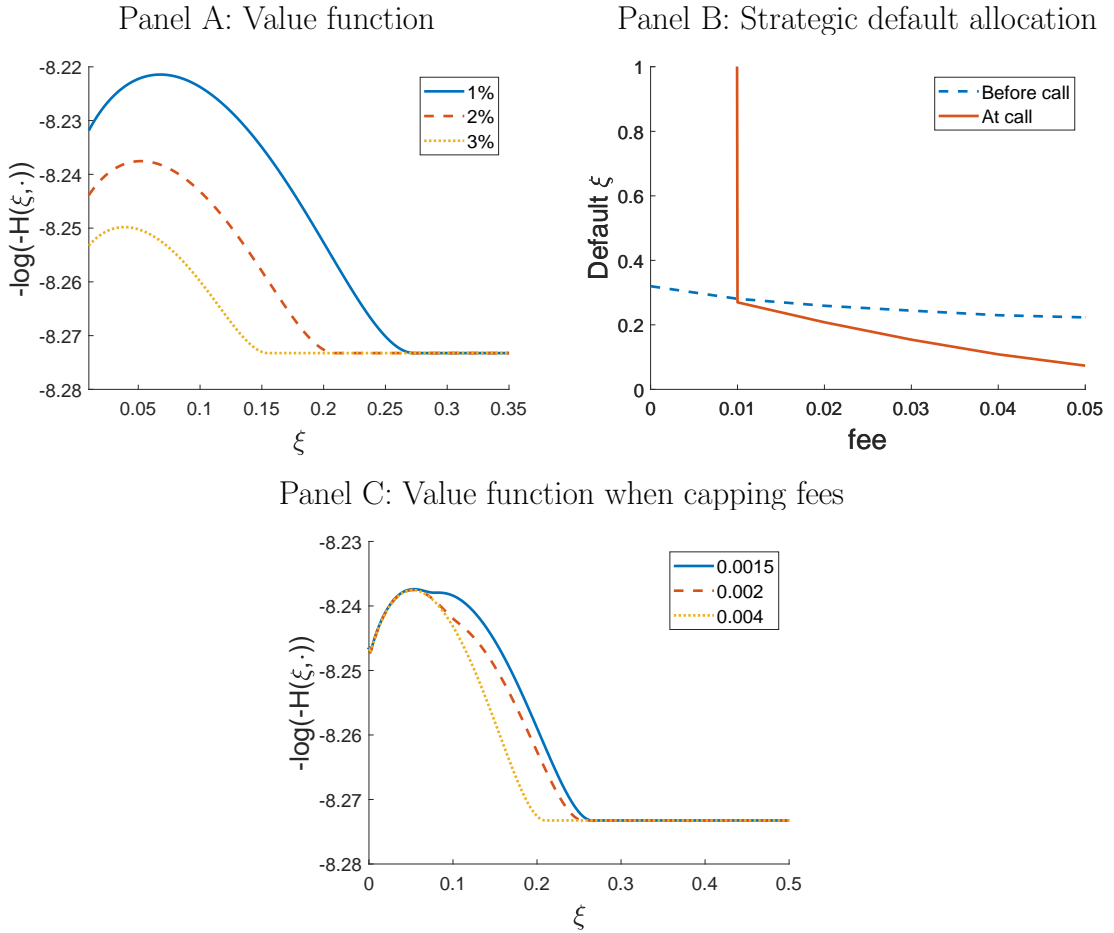
Figure 5. Sensitivity of Optimal PE Allocation to Model Parameters

This figure shows the optimal PE commitment ξ^* as we vary model parameters. Each Panel A to H varies one parameter around the standard calibration value (dotted vertical line), keeping the others fixed at their values in our standard calibration (Table 1). The solid line represents PE commitment in the baseline economy (E0); the dashed line is for the economy without commitment-quantity risk (choose quantity on call, E2); the dot-dash red line is for an economy in which PE is fully liquid (the Merton two-risky-asset model).

Panel A: PE returns ν

Panel B: PE volatility ψ

Panel C: Stock returns μ

Panel D: Stock volatility σ

Panel E: Correlation between PE and stock ρ

Panel F: Risk aversion γ

67
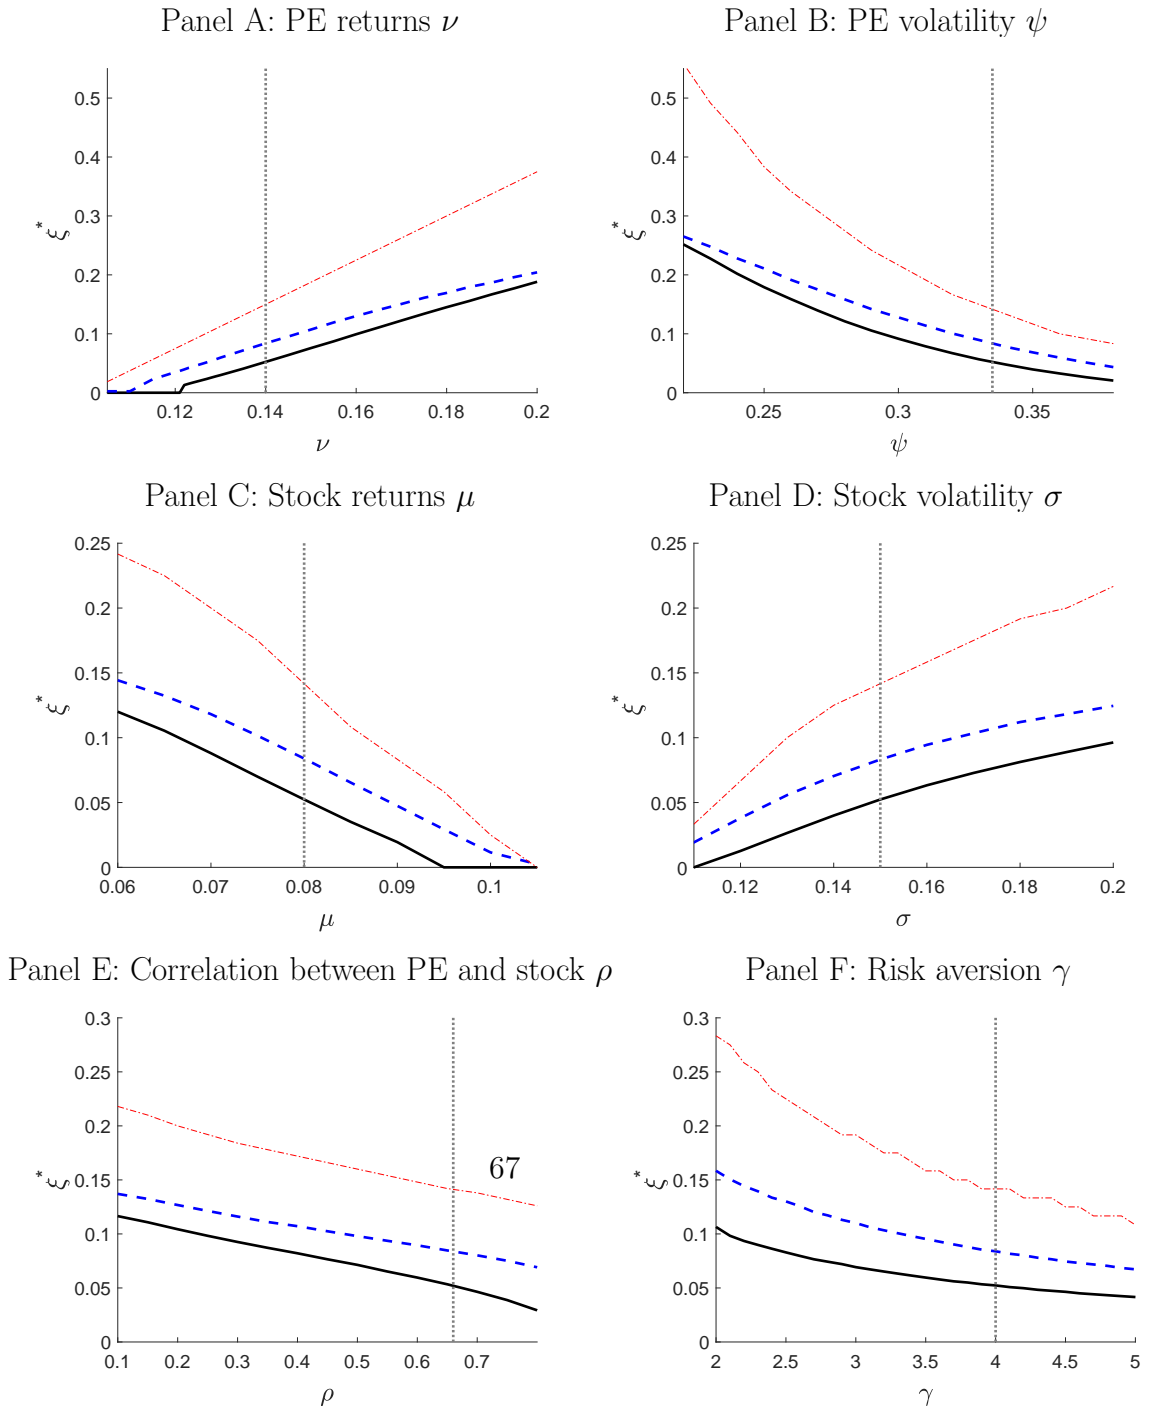
Figure 6. Sensitivity of Return Premium and Welfare Cost of Commitment-Quantity Risk to Model Parameters

This figure represents the return premium and the welfare cost of commitment-quantity risk, as functions of the risk-return parameters of the model. We only vary one parameter while keeping the other parameters constant. The fixed parameters are those of our standard calibration (Table 1). The vertical dotted lines represent the value used in the standard calibration for the parameter that we vary.

Panel C: Stock returns μ

Panel D: Stock volatility σ

Panel E: Correlation between PE and stock ρ

Panel F: Risk aversion γ
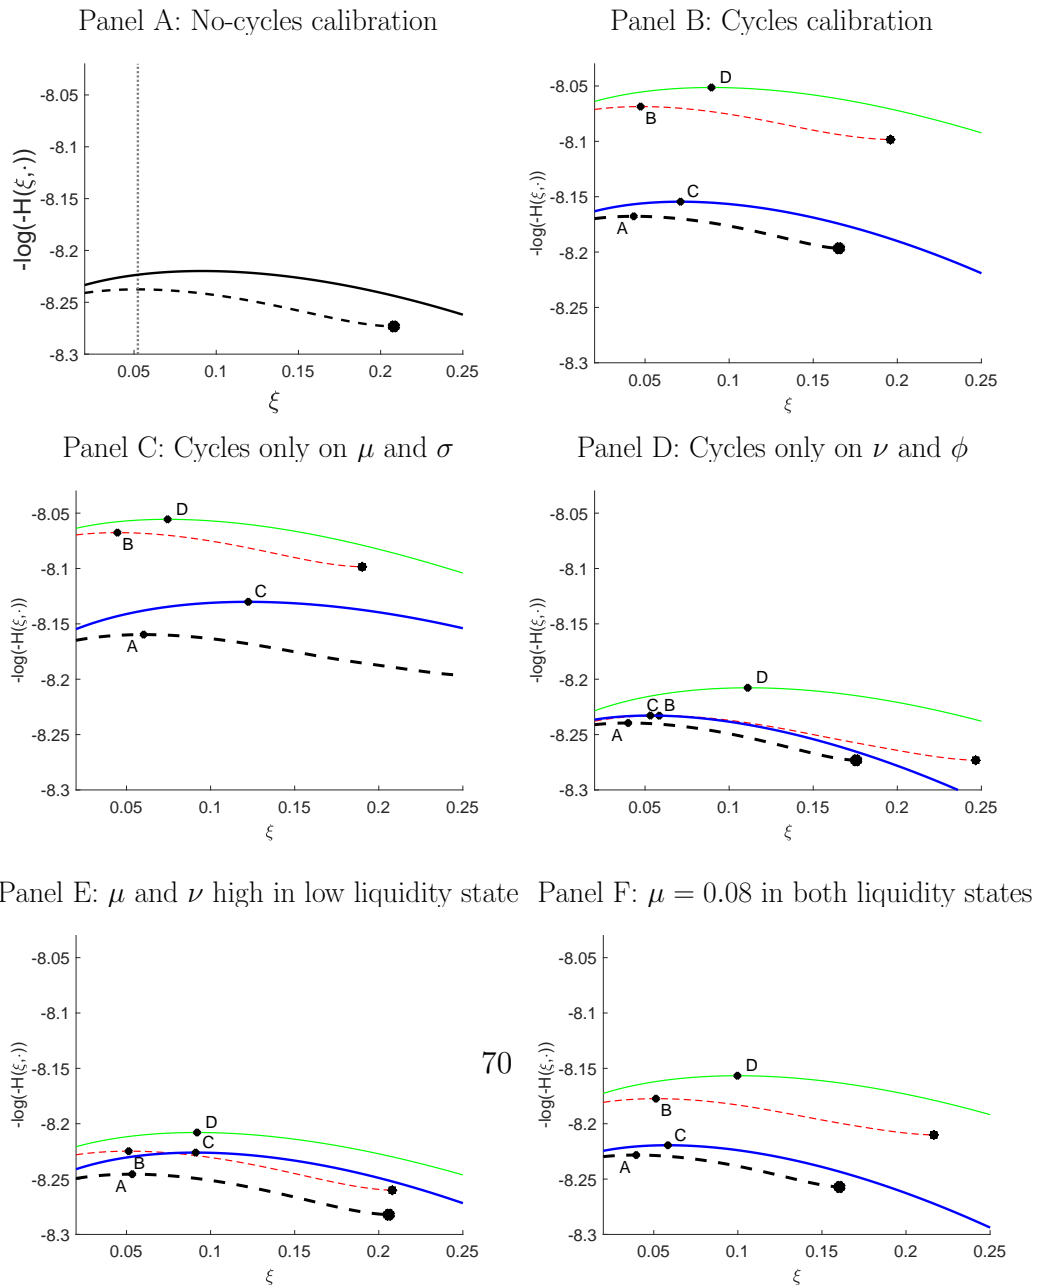
Panel G: Intensity of calls λ_C

Panel H: Intensity of distributions λ_D

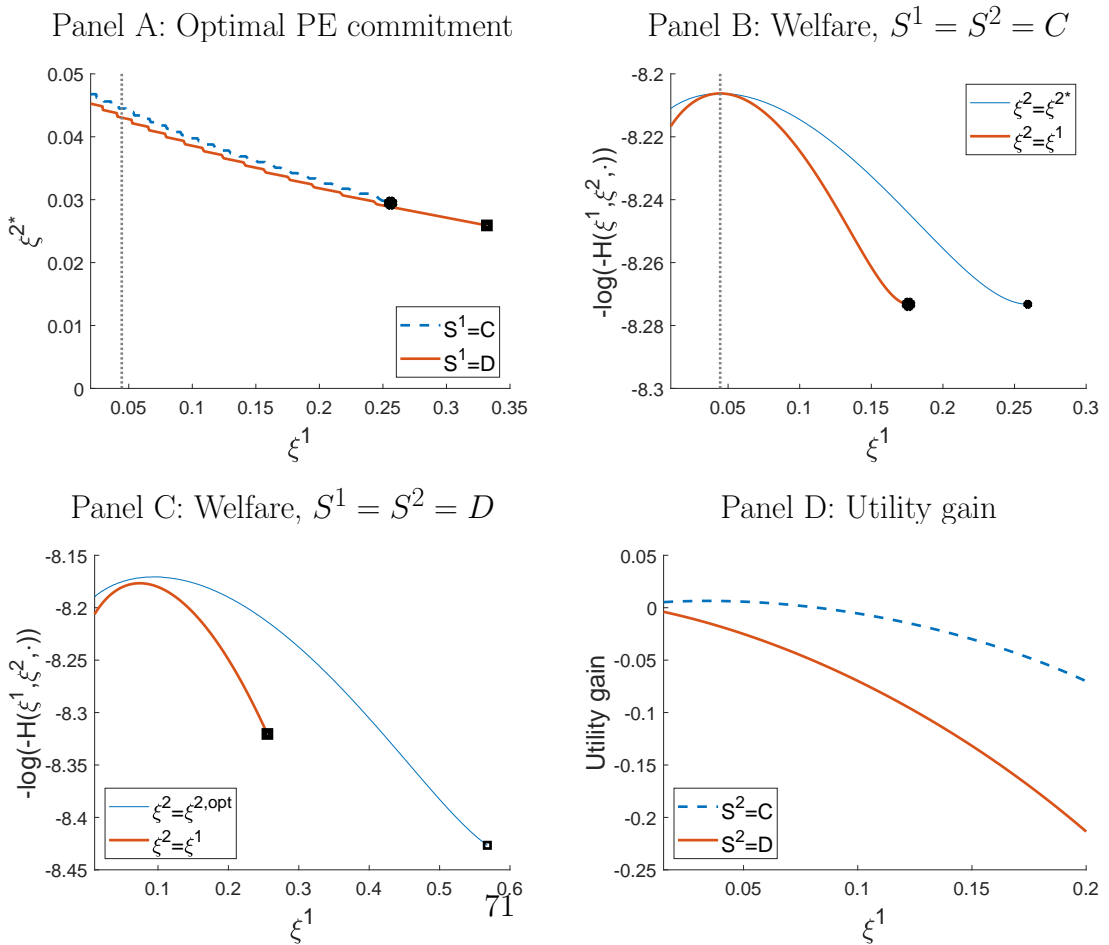**Figure 7.** Welfare in the Baseline Economy with Cycles

This figure represents the investor's value function in the low liquidity state during the commitment period (thick dashed line, point A marks the optimal PE commitment) and the holding period (thick plain line, point C marks the optimal PE allocation), and in the high liquidity state (thin dashed line and thin plain line, B and D mark the optimal commitment and allocation). Panel A uses the calibrated cycles parameters (Table 5). In the other panels, we keep all calibrated parameters except those indicated. In Panel B, μ is set to its value in the economy with no cycles ($\mu = 8\%$ in both liquidity states). In Panel C, μ and ν are high (low) in the low (high) liquidity state: $\mu = 8.6\%$ and $\nu = 15.3\%$ ($\mu = 7.7\%$ and $\nu = 13.4\%$). In Panel D, μ is high (low) in the low (high) liquidity state: $\mu = 8.6\%$ ($\mu = 7.7\%$). In Panel E, ν is high (low) in the low (high) liquidity state: $\nu = 15.3\%$ ($\nu = 13.4\%$).

Panel A: No-cycles calibration

Panel B: Cycles calibration

Panel C: Cycles only on μ and σ

Panel D: Cycles only on ν and ϕ

Panel E: μ and ν high in low liquidity state

Panel F: $\mu = 0.08$ in both liquidity states

70

Figure 8. PE Allocation and Welfare in the Two-Fund Model

Panel A represents the optimal commitment of the investor to the second fund, ξ^{2*}, as a function of their ongoing allocation to the first fund ξ^1, if the first fund is (i) in the commitment period (dashed line) and (ii) in the holding period (solid line). The circle indicates strategic default in fund 1. The square indicates fund 1 being sold on the secondary market. The dotted vertical line marks the optimal allocation in the first fund, ξ^{1*}. Panels B and C display the value function of the investor as a function of the allocation in fund 1, (B) when both funds are in the commitment period, (C) when both funds are in the holding period. Thin lines correspond to the case in which the allocation in fund 2 is optimal (B) at inception of the fund and (C) at capital call. Thick lines correspond to the case in which the same fraction of wealth is allocated to both funds. Panel D represents the utility gain when fund 1 calls, if fund 2 is in its commitment period (dashed line) and if it is in its holding period (solid line). The utility gain is defined as the difference in the log value function after adjusting for the jump in liquid wealth at capital call.

Panel A: Optimal PE commitment

Panel B: Welfare, $S^1 = S^2 = C$

Panel C: Welfare, $S^1 = S^2 = D$

Panel D: Utility gain

71

Figure 9. PE Allocation in the Infinite-Fund Case

This figure represents the optimal aggregate commitment of the investor, π^*, as a function of the capital that is already invested, ξ, in the infinite-fund model. The diagonal line represents the committed capital as a function as the invested capital in the steady state, see equation (17). The steady state is defined as the state in which the expected change in invested capital is zero: $E[dY_t^\infty] = 0$, where Y_t^∞ follows the dynamics given in equation (12). The dotted vertical line marks the optimal investment in the steady state.

Figure 10. Welfare Cost of Timing Risk

This figure represents the welfare cost of commitment-timing risk as a function of the subjective discount factor δ. We only vary δ while keeping the other parameters constant and as given in our standard calibration (Table 1). In the standard calibration we use $\delta = 0.05$ (vertical dotted line). The horizontal line is at zero.