Capital Commitment *

Elise Gourier
ESSEC Business School and CEPR

Ludovic Phalippou
University of Oxford, Said Business School

Mark M. Westerfield
University of Washington, Foster Business School

April 2021

Abstract

Ten trillion dollars are allocated to illiquid vehicles for which investors commit ex-ante to transferring capital on demand – most of which are Private Equity (PE) funds. We design a dynamic portfolio allocation model in which investors commit capital to PE. Investors significantly under-commit and are willing to pay as much as 15% of their PE allocation to change the amount committed. A more liquid secondary market or access to multiple PE funds further *increase* the required compensation for commitment risk. The uncertainty about the timing of capital calls and the penalty in case of default have minor effects.

JEL Classification: G11, G12

Keywords: Capital commitment, private equity, liquidity cost

---

*elise.gourier@essec.edu, ludovic.phalippou@sbs.ox.ac.uk and mwesterf@uw.edu. We are thankful to Andrew Ang, Michael Brandt, John Campbell, Joost Driessen, Lubos Pastor, Jos van Bommel, Marno Verbeek, and Luis Viceira for early comments and encouragements on our project. We also thank participants to the Luxembourg Asset Pricing and Amsterdam Netspar conferences for comments. This research has been supported by a Netspar grant and a grant from the "Institut Europlace de Finance".

1 Introduction

Institutional investors' allocation to private market funds nears the $10 trillion landmark.[1] These funds span a wide range of investments from real estate to leveraged buyouts and venture capital. A unique feature of private market funds, irrespective of their focus, is that they require investors to commit capital ex-ante to fund managers and, therefore, to relinquish control on investment timing and quantity. This paper is the first to study the effect of capital commitment on portfolio allocation decisions and investors' welfare.

Consider the asset allocation of Yale University Endowment: 26% hedge funds, 24% liquid assets, and 50% private equity funds. Their capital commitments stand at 33% of their assets under management, which is more than the value of their liquid asset holdings.[2] The 2008 financial crisis showed that large capital commitments can impose a significant cost for institutional investors. Pension funds (e.g., Calpers), universities (e.g., Harvard), and foundations (e.g. Carnegie) attempted to rebalance their portfolios away from private market funds at great costs.[3] Since 2008, capital commitments have doubled. Globally, capital commitments (a.k.a. dry powder) now add up to as much as $3 trillion. However, in the meanwhile, a secondary market for private market funds has emerged.

In this paper, 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 find that the effects of commitment on investors' welfare and portfolio are large. In particular, they are much larger than investors' welfare gain from having a liquid secondary market.

We model investments in private equity funds as follows. At time 0, the investor commits a positive amount to private equity and does not know either when her capital will be called, or when investment proceeds will be distributed. We use a Poisson process to model the stochastic timing of capital calls and payouts. The first jump triggers the capital call, date at which the investor transfers the committed amount to the fund manager. If the investor is unable to make this transfer, she defaults on her pledge. The second jump of the Poisson process marks the time at which the fund manager distributes the proceeds from the investment back to the investor. At the same time, a capital commitment is made to a new private equity fund, and the above is repeated all the way to infinity.

---

[1] Information about assets under management and dry powder is from the Preqin dataset; end of year 2020.

[2] https://your.yale.edu/sites/default/files/20192020-annual-financial-report.pdf

[3] Institutional investors resorted to costly strategies such as redeeming capital from other investments despite low overall liquidity, selling private market funds on the secondary market at large discounts, and issuing high yield bonds. See Barron's, 69009,"The Big Squeeze"; Forbes, 104009, "Did Harvard Sell At the Bottom?"; Institutional Investor, 11009, "Lessons Learned: Colleges Lose Billions in Endowments."

We define and isolate the different types of liquidity frictions. Two frictions stem from not being able to immediately deploy the capital into the illiquid asset (commitment risk): i) the timing of the capital call is unknown; ii) the amount of capital that will be called *as a fraction of wealth* is unknown.[4] We define the welfare cost of *commitment-timing* risk to the investor as the amount she is willing to pay to know when the capital call will occur. Similarly, the cost of *commitment-quantity* risk is the amount the investor is willing to pay to update her commitment at the time of capital call. In addition to these frictions linked to commitment risk, additional illiquidity stems from the non-tradability of the private equity fund, as studied in Ang et al. (2014).

This setup is sufficiently flexible to incorporate three realistic features of private markets: i) strategic default: the investor can skip a capital call, at the cost of no longer being able to invest in private equity; ii) access to a secondary market: the investor can sell her invested private equity holdings at a discount; iii) contemporary commitments to multiple private equity funds. We can, therefore, not only analyze the different frictions induced by capital commitment, observe whether some are more important than others, but also study the interactions between these frictions and different features of private markets.

We conduct a unique and thorough calibration of our model. We search for the combination 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 obtain a correlation between private and public equity of 80%. The drift and volatility are 14.5% and 25% for private equity returns. These returns are, however, gross of management fees paid during the commitment period. The (net-of-fees) public equity returns drift and volatility are calibrated at 8% and 15%. These parameters imply a private equity beta of 1.32, which coincides what has been found in the literature using different approaches (e.g. Ang et al. (2018)).

The expected PME distribution that results from our calibration matches the empirical one remarkably well, despite a weakness at reproducing the relatively high frequency of low PMEs. As a result, our optimal commitment to PE is quite high, at 22% of wealth.[5]

---

[4] During the time period between commitment and capital call (commitment period), the investor's wealth faces random shocks, causing the committed amount to be sub-optimal at the time of the capital call. Following a decline in stock prices during the commitment period, the amount called may for example be larger than the optimal amount.

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

Our findings are as follows:

- **Large under-commitment to PE**. On average, investors significantly under-commit to private equity from an ex-post perspective. This result rationalizes the increased offering of co-investment opportunities to investors by fund managers at the time of capital call. We show that co-investment opportunities are valuable options to increase exposures ex-post, whereas the literature mostly presents them as tools to reduce fees.

- **The cost of commitment-quantity risk is large**. Whenever capital is called to be invested, the asset allocation is sub-optimal because the stock-market has moved since the time of commitment. Investors are willing to pay 5% of their wealth (15% of their PE allocation) to be able to change their PE allocation when capital is called.

- **Other commitment risks have minor effects**. The cost of commitment-timing risk is small and slightly *negative*. Investors receive a utility gain at capital call, whose value is exponentially discounted. Due to the convexity of exponential discounting, they prefer uncertainty to certainty over the timing of this gain.

- **Liquid secondary markets are not valuable per se**. A more liquid secondary market for PE funds leads to a modest increase in the optimal PE allocation.[6]

- **Secondary market liquidity and commitment risk are complements**. Liquid secondary markets exacerbate the cost of commitment-quantity risk: investors are willing to pay even more (6% of their wealth) to be offered the possibility to change their PE allocation when capital is being called if the secondary market is more liquid.

- **No time diversification benefits**. With access to multiple funds, the cost of commitment-quantity risk further increases to 6.4% of wealth. This result is the opposite of a naive prediction based on time diversification but is consistent with different types of liquidity being complements. When the commitment on one of the funds has been called, the investor does not know whether she will receive the proceeds of the investment before the commitment on the second fund is called. In the event that the second fund calls before the first fund distributes, the investor faces a liquidity mismatch. Also, following a negative shock to stock prices, the commitment on the second fund may become large relative to liquid wealth, which increases default risk. She then decreases her allocation to stocks and increases that to bonds.

[6]Our secondary market is restricted to the trading of private equity funds that have called their capital. We do not allow for unfunded commitments to be traded. In practice, secondary markets are most active for funds that are fully invested (Nadauld et al. (2019)).

2 Institutional Setup & Literature Review

2.1 Fund Structure

Private Equity (PE) funds are set up as private limited partnerships and pursue a wide range of investment strategies (e.g. leveraged buyout, real estate, venture capital, infrastructure, mezzanine etc.). A PE firm acts as a General Partner (GP) in the limited partnership and the capital providers are the Limited Partners.

During a fund-raising period that spans three to eighteen months, the PE firm (e.g., KKR) seeks capital for its fund (e.g., KKR millennium fund). Investors admitted as Limited Partners (LPs) commit a certain amount of capital, i.e. agree to provide cash to the fund on demand up to that committed amount. When the PE firm ends its fundraising, it has its 'final close' and the year this occurs is called the fund vintage year. The total capital committed to a fund is referred to as the fund size.

Capital calls are usually made in connection to either a fee payment or a specific investment, but some pooling occurs to reduce the frequency of capital calls to Limited Partners. This pooling of capital calls is enabled by credit facilities contracted by the General Partner using LP commitments as collateral. It is also the case that institutional investors are exposed to private equity funds via intermediary vehicles (e.g. fund-of-funds, separately managed accounts) which may also partially pool capital calls.

Both the timing and the size of the capital calls (also known as draw-downs) are uncertain for LPs. They only know an ex-ante specified investment period, during which most capital calls should occur. For leveraged buyout funds, for example, the investment period is typically five years. Some capital may be called after this investment period for fee payments or follow-on investments in existing portfolio companies. For venture capital funds, the investment period is longer than for buyout funds to allow for large follow-on investments in portfolio companies that are successful.

When an investment is exited, the payout is usually distributed to LPs and cannot be recycled to make a new investment, but there are some exceptions. The divestment period is flexible, spanning the entire life of the fund, including an overlap with the investment period. Funds life is 10 years but there are multiple circumstances under which funds obtain extensions. We observe many funds that are not fully liquidated in their fifteenth year and beyond.

2.2 Defaulting on Capital Calls & Secondary Markets

LPs defaulting on capital calls is a rare event. Penalties for doing so can be high: forfeiture of some or all existing investments in the fund, and impossibility to invest in PE going forward. The 2008 financial crisis shows how LPs incurred a cost due to their unfunded commitments: "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."[7]

But, in the end, they all avoided default on PE commitments. In addition, once incipient due to contractual restrictions on transfers (see Lerner and Schoar (2004)), the secondary market for PE fund stakes took off right after the 2008 financial crisis to an annual turnover of $88 billion in 2019.[8] Specialized secondary funds raised as much as $100 billion in 2017-2020 to buy fund stakes.[9] The increasing volume of trades suggests that PE funds are much less illiquid than they used to be. Applying the practitioners' judgement, which measures illiquidity by the observed discount to the reported Net Asset Value (NAV) in each transaction, PE investing is quasi-liquid because the average discount since 2010 has been less than 10% and zero in several quarters. However, and importantly, liquidity varies greatly across fund types and fund age. For example, there are few transactions of funds during their investment period, and most transactions are for relatively large funds (Nadauld et al. (2019))

[7] From the magazine 'Private Equity Insider' in its issue of November 5, 2008.

[8] Volume passed $10 billion in 2006, $20 billion in 2010, $40 billion in 2014.

[9] Other buyers include funds-of-funds, which raise third party capital to invest in both the primary and secondary market; and various asset owners (investment banks, hedge funds, sovereign funds), which have dedicated teams to secondary market purchases.

2.3 Co-Investments

Another major transformation of the structure of the PE industry post crisis is the growth in LP co-investments and direct investments. In a nutshell, when GPs call capital for an investment, they select some LPs to whom they offer the possibility to co-invest, i.e., add additional capital in that deal. Direct investment refers to a situation where an LP directly executes a PE transaction: e.g., a pension fund buys an operator of nursing homes.

Pre-2008, co-investment invitations were limited to large global LPs and to other PE firms in so-called club deals (Officer et al. (2010)). Post-2008, invitations are more widespread but mostly restricted to LPs in the PE fund that makes the transaction. These co-investments remain a small fraction of the total amount committed to PE but their prevalence is growing fast.[10] Similarly, direct investments remain a small fraction of overall PE investments, but are also fast-growing.

The main motivation LPs give for co-investments is to lower fees (Da Rin and Phalippou (2017), Fang et al. (2015)). A second motivation – highlighted by Braun et al. (2020) is that: "co-investments give the opportunity for LPs to increase their exposure, above the pro rata interest obtained via their fund investment, in particular companies or sectors. This could be driven by a belief that investors can spot particularly attractive portfolio companies. Alternatively, investors could use co-investment to tilt their portfolio toward particular sectors, such as new technologies, that are difficult to gain exposure to via public markets. While co-investments allow investors greater control, they need to have the relevant skills and experience to evaluate which transactions to participate in."

The above comment is particularly important in the context of our model. It means that, in practice, LPs have the possibility to adjust upwards, but not downwards, their allocation to PE. Direct investments offer a similar opportunity. Again, the amounts are still modest but they are growing fast and this is happening at the same time that both secondary markets and allocation to PE funds have grown significantly. We also note that Fang et al. (2015) find some evidence that when LPs find themselves under-exposed to PE, they accept more co-investment invitations, possibly to ramp-up their PE allocation. Similarly, it is routinely argued that some LPs buy funds on the secondary market to be more quickly exposed to PE than if they go with the usual route of committing capital to newly raised funds.

[10] Triago, a PE specialist, estimates that in 2019, \$66 billion was co-invested by LPs. For comparison, amount committed to PE funds is about \$1 trillion yearly. Lerner et al. (2020) show that in 2017, 40% of all capital raised in PE was done via alternative vehicles (a.k.a. 'shadow capital'), but this statistic also include things like feeder funds and separately managed accounts.

2.4 Literature Review

Several papers look at the investor's portfolio choice problem in the presence of an illiquid asset, defined as an asset that the investor cannot trade during a given period of time. In Dai et al. (2015), this period is deterministic. In Ang et al. (2014), like in our setting, this period is stochastic. In Dimmock et al. (2019), this period is deterministic but they allow for a secondary market: investors can voluntarily transact in the illiquid asset at any time by paying a proportional transaction cost. They also allow for multiple illiquid assets. In Giommetti and Sorensen (2020), funds are gradually called and distributed from a composite private equity fund.

Our model allows for more than one illiquid asset, solves an infinite horizon problem, has a secondary market, has periods of illiquidity that are stochastic, and incorporates capital commitments.

In a related setup, Sorensen et al. (2014) model illiquidity by constraining an investor to hold a private equity fund until maturity. The private equity fund is acquired at time 0, hence the capital is immediately put at work by the fund. There is therefore no capital commitment or possible default on the committed capital. The fund is liquidated at maturity T, which is finite and known ex ante. The main objective of the paper is to extend the valuation of the performance-related fee (carry) charged by general partners solved in Metrick and Yasuda (2010) to a setting where funds are not continuously traded.

Longstaff (2001, 2009) is the first to model how the presence of an illiquid asset affects optimal portfolio decisions and pricing. The model is in general equilibrium and features a representative agent endowed with an illiquid asset. This asset becomes liquid at a random time. The problem has finite horizon, with maturity when the illiquid asset pays off.

There is also a small empirical literature on private equity illiquidity. Ljungqvist et al. (2017) propose a hazard/duration regression approach to model the speed at which capital is called as a function of market conditions for a sample of US buyout funds. In a similar exercise, Robinson and Sensoy (2016) find that fund net cash flows are pro cyclical, i.e. distributions tend to decrease and capital calls tend to increase in bad times, suggesting high liquidity risk. However, they argue that this liquidity might be diversifiable. We show that some liquidity risks seem benign indeed but commitment risk is not, and that commitment risk is not diversifiable. To the contrary, the presence of more funds in a portfolio increases the probability of a liquidity mismatch.[11]

[11] Also related is Franzoni et al. (2012) who empirically quantify the liquidity risk premium for private equity in the four factor model of Pastor and Stambaugh (2003). They also show that private equity returns are particularly sensitive to changes in aggregate liquidity conditions during the life of the investments.

3 Model

We consider an information structure that obeys standard assumptions. There exists a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ supporting the vector of two independent Brownian motions $Z_t = (Z_t^1, Z_t^2)$ and an independent Poisson processes (M_t). $\mathcal{P}$ is the corresponding measure and $\mathcal{F}$ is a right-continuous increasing filtration generated by $Z \times M$.

3.1 The Three Assets

There are two liquid assets in the economy in the sense that they can be rebalanced continuously at no cost: i) a risk-free asset, which captures the fixed-income market, labelled bond; and ii) a risky asset, which captures the public equity market, labelled stock.

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

$$dB_t = rB_t dt \tag{1}$$

The stock price P_t follows a geometric Brownian motion:

$$\frac{dP_t}{P_t} = \mu\, dt + \sigma dZ_t^1 \tag{2}$$

where Z_t^1 is a standard Brownian motion, μ is the return drift, and σ the return volatility.

Investors can allocate capital to a third asset: a private equity fund. As illustrated below, the fund manager collects capital commitments to this third asset from the investor at time $t = 0$, then calls the committed capital at time $t = \tau_C$ and pays out the value of the investment at time $t = \tau_D$. At this point the next commitment is being made and the above is repeated, all the way to infinity. We will call the *commitment period* the time period $[\tau_0, \tau_C]$, and the *holding period* the time period $[\tau_C, \tau_D]$.

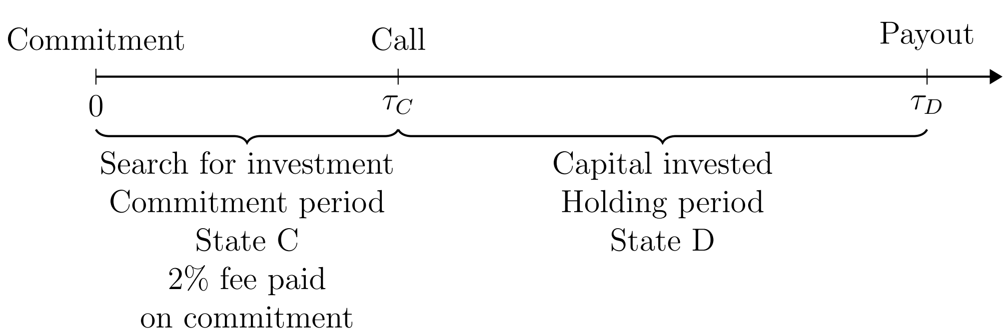

3.2 Modelling Private Equity

At time τ_0, the investor commits a positive amount X_0 to the private equity fund. This commitment is a promise to make capital available when the manager calls it at time $\tau_C > 0$. During the commitment period $[\tau_0, \tau_C]$, fees are being paid to the fund manager (set to $0.02 \times X_0$), but no investment is made.

In a slight abuse of notation, we denote by X_t the amount of capital committed between times τ_0 and τ_C, and use X_t again to refer to the amount of capital invested in the fund after τ_C. Because no capital is invested at time 0 and the commitment cannot be changed, $dX_t = 0$ until τ_C (time at which the committed amount is called.)

We use a Poisson process M_t to model the transfers of capital between the investor and the fund manager. The process starts at $M_0 = 0$ and has initial intensity λ_C. The first jump triggers the capital call, at $t = \tau_C$, date at which the investor transfers X_0 of liquid wealth to the fund manager.

At time τ_C, the intensity of the Poisson process changes and becomes λ_D, and the capital committed is invested in the private equity fund at that point in time.[12] During the holding period $[\tau_C, \tau_D]$, the value of the private equity investment, net of all fees evolves as a geometric Brownian Motion, with drift ν, and volatility ψ:

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

Where $dZ_t = \rho dZ_t^1 + \sqrt{1-\rho_L^2} dZ_t^2$

This setup implies that the correlation between public and private equity is ρ; and

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

The Poisson process M_t jumps at time $t = \tau_D$, at which point the private equity investment is fully exited, and the investors receive the value of the private equity fund: X_{τ_D}. The process M_t is reset to zero, and the investor immediately makes a new capital commitment X_t to a new private equity fund. The intensity of the process M_t switches back to λ_C, and the above is repeated.

[12]In terms of fees, we can either switch the fee base from capital committed to asset value, or ignore fees after commitment. These two options are strictly equivalent, the former simply makes notations heavier, and we thus opt for the latter. What we can not easily do is switch the fee base to invested capital because that would require another state variable.

During the holding period, the investor can sell her private equity stake on a secondary market. In this case, she receives αX_t. At that point, she makes a new commitment to private equity and the above is repeated (wait until capital is called, and then until it is paid out). If the investor is unable to pay for the capital call at time τ_C, i.e. $X_0 \geq W_{\tau_C}$, she defaults on her commitment. The consequences of such a default is that she is forever banned from accessing private equity – which is a realistic feature (Banal-Estañol et al. (2017)).[13]

To ensure parsimony and analytical tractability, we made several simplifying assumptions. First, as explained in Section 2, in practice, capital calls are spread across the investment period, with the sum of the capital calls equal to the committed amount. In our model, there is a single capital call equal to the committed amount. Similarly, our model has a single payout. This stylized representation allows us to have state equations that have no memory, and to use numerical methods based on Markov chain approximations to solve the portfolio allocation problem.

Second, the investment opportunity set contains a single private equity fund at any point in time. We extend our model to a two-fund setup in Section 6.4. The model can actually be extended to any number N of funds. However, the computational cost becomes quickly prohibitive. Each fund can either be in state $S = C$ or $S = D$, the number of state variables thus equals 2^N.

Third, our fee structure is highly simplified. Our most important omission is that of the carried interest. Since the carried interest calculation requires the knowledge of X_0 at time $t = \tau_D$, we would need an additional state variable. Similarly, the base of the management fees during the holding period is usually the amount invested, which is also X_0 in our model. Hence, choosing this fee base would also require to track X_0, and thus an additional state variable. However, given the complex nature of the capital structure used in private equity transactions and the multiple operational changes private equity funds implement, it is an open empirical question as to whether it is the net of fees or the gross of fees private equity value dynamics that is closest to a geometric Brownian motion.

[13] Note that the investor can only sell her full investment in a fund on the secondary, which is realistic. She cannot sell her commitment on the secondary market, however, which is also realistic as such transactions are extremely rare.

3.3 The Investor Problem

At each time t, the investor chooses how to allocate her liquid wealth between the two liquid assets. θ_t denotes the fraction of liquid wealth allocated to stocks at time t.[14] In addition, at time τ_0, and each time a private equity fund is liquidated, she chooses how much capital to commit to private equity (X). Finally, she decides her consumption rate (out of liquid wealth), $c_t = C_t/W_t$.

Hence, the evolution of the investor's liquid wealth W_t, is given by:

$$\frac{dW_t}{W_t} = (r + (\mu - r)\,\theta_t - c_t)\,dt + \theta_t \sigma dZ_t^L - \frac{dI_t}{W_t} \tag{5}$$

where dI_t denotes any transfer between illiquid wealth and liquid wealth. The optimal value function is then 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 constraints (3) and (5), with δ denoting the subjective discount factor, and $S_t = \{C, D\}$ denoting the state.

When $t = \tau_C$, $dI_{\tau_C} = X_0$; i.e. committed capital is called, and dI_t is transferred out of liquid wealth to private equity. The value function jumps discretely from $F(W, X, S)$ to $F^{+C}(W, X, S) = F(W - X, X, S)$. The state S changes from C (commitment period) to D (holding period).[15]

When $t = \tau_D$, $dI_{\tau_D} = -X_{\tau_D}$; i.e. capital is paid out and dI_t is added to liquid wealth. The value function jumps discretely to $F^{+D}(W, X, S) = F(W + X, X, S)$. The state changes from D to C.

When $t \neq \tau_C$ and $t \neq \tau_D$, $dI_{\tau_D} = -X_{\tau_D}$, and $dI_t = 0$.

In the base case, the investor has a standard power utility U, i.e. has Constant Relative Risk Aversion (CRRA) over sequences of consumption, $U(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$, with $\gamma > 1$.

[14]Following Dybvig and Huang (1988) and Cox and Huang (1989), we restrict the set of admissible trading strategies, θ, to those that satisfy the standard integrability conditions. All policies are appropriately adapted to $\mathcal{F}_t$.

[15]We drop the time subscript for notation simplicity.

3.4 Solution

To avoid the objective function to be infinite, optimal policies are such that the liquid wealth W is at all time positive: $W > 0$ and the illiquid wealth is bounded from below by zero: $X \geq 0$, almost surely (a.s.) At the time of capital commitment, the amount X is chosen so that $X < W$ a.s. at any time in the future. If $X \geq W$ at the time of the capital call, the investor defaults on her commitment, driving her utility to infinity.

The investor's optimal consumption and allocation in the three assets solve the Hamilton-Jacobi-Bellman (HJB) equation:

$$\beta F = \max_{\{\theta, X, c\}} \Big\{ U(c) + (r + (\mu - r)\theta - c)\, W F_W + \frac{1}{2}\theta^2\sigma^2 W^2 F_{WW} + \nu X F_X + \frac{1}{2}\psi^2 (X)^2 F_{XX} \\ + \theta\sigma\psi\rho_L W X F_{WX} + \lambda_C \left(F^{+C} - F\right) \mathbb{1}_{S=C} + \lambda_D \left(F^{+D} - F\right) \mathbb{1}_{S=D} \Big\}. \tag{7}$$

Because the utility function is homothetic and the return processes have constant moments, it must be the case that the value function F is homogenous of degree $1-\gamma$ in total wealth. If we denote TW the investor's total wealth, i.e. the sum of liquid and illiquid wealth:

$$TW \equiv W + X\mathbb{1}_{S=D} \tag{8}$$

where the indicator $\mathbb{1}$ is one if state id D, then the value function can thus be written as the product of a power function of total wealth and a function of the wealth composition:

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

where $\xi = \frac{X}{TW}$ denotes the ratio of private equity to total wealth and is required to be smaller than one in state D.

Proposition 1 *The investor's value function can be written as in (9), where $H(\xi, S)$ exists and is finite, continuous, and concave for $\xi \in (1, \infty)$. Whenever the investor can commit capital, she selects $\xi^* \equiv \arg\max_\xi H(\xi, S = C)$, which exists.*

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

3.5 The Illiquidity Stack

The model (3)-(6) is referred to as the **baseline economy (Economy 0)**, i.e. an economy where all the sources of private equity liquidity risk are present. We now define seven other economies. Each economy corresponds to a situation where one or several of the features of our baseline model are modified.

In the baseline economy, the investor does not know *when* she will need to pay the committed capital out of her liquid wealth. To turn this *commitment-timing risk* off, we make the call time deterministic.

Furthermore, in the baseline economy, the investor does not know *how much* she will need to pay into private equity relative to her liquid wealth. As her wealth evolves randomly between commitment time and call time, this relative amount is random and no longer optimal at call time. 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. In this case, the investor is no longer exposed to default risk.

If neither commitment-timing nor commitment-quantity risk is present, there is no commitment risk. Absent commitment risk, our model still features a delayed investment into private equity; the investor still needs to wait until call time to access the private equity returns. This restriction is lifted in economy 4.

Economies 5 to 7 below still have commitment risk, and focus on the liquidity frictions during the holding period.

Economy 1: Deterministic call time, choose quantity when committing

The length of the commitment period is set to $\frac{1}{\lambda_C}$, *commitment-timing risk* is therefore switched off. The investor is still exposed to commitment-quantity risk: the *amount* called relative to liquid wealth at call time τ_C is uncertain.

Economy 2: Stochastic call time, choose quantity when called

In this economy, *commitment-quantity risk* is turned off. The commitment-to-wealth ratio ξ^* is no longer chosen at commitment time $t = 0$, but at call time (τ_C). The timing of the capital call remains stochastic.

Economy 3: Deterministic call time, choose quantity when called

Both *commitment-timing risk* and *commitment-quantity risk* are turned off. This economy is one without *commitment risk*. The private equity fund remains illiquid and the payout time remains uncertain.

Economy 4: Immediate private equity access, no commitment risk

The commitment period time is brought to zero. Hence, at time $t = 0$, and after each private equity fund liquidation, the investor freely allocates capital between bonds, stocks and private equity. Private equity remains illiquid: capital allocated to private equity is unavailable until a stochastic payout time τ_D (unless the investor opts to sell on the secondary market).

Economy 5: Deterministic payout time

To turn off the payout-timing risk, the capital distribution is no longer random. Instead, it is deterministic and equal to $\frac{1}{\lambda_D}$. Payout-timing risk was introduced under the label 'illiquidity risk' by Ang et al. (2014) – it refers to the uncertainty about when the illiquid asset will pay out.[16] This economy still contains commitment risk.

Economy 6: Fully liquid secondary market, with commitment risk

During the commitment period, the investor can sell her private equity stakes on the secondary market at no cost. This economy still contains commitment risk.

Economy 7: Secondary market with 40% discount

During the commitment period, the investor can sell her private equity stakes on the secondary market at a 40% haircut. This economy still contains commitment risk.

We define the **welfare cost to the investor** of any economy A with respect to any economy B as the fraction of wealth ζ the investor would be willing to pay to switch from economy A to economy B:

$$H^A(\xi^{A*}, S = C) = (1 - \zeta)^{1-\gamma} H^B(\xi^{B*}, S = C), \tag{10}$$

where ξ^{A*} and ξ^{B*} denote, respectively, the optimal commitments in economies A and B. If the cost ζ is zero, then the investor is indifferent between the two economies. We measure these costs at time $t = 0$ (right before the investor makes a capital commitment).

[16]Note that Longstaff (2009) introduces a general equilibrium model with an illiquid asset, which cannot be traded between inception and a deterministic time T. After this time T, the asset becomes liquid again.

4 Model Calibration

We restrict the investment opportunity set to the US and to three assets (a risk-free asset, public equity and private equity), and use the past thirty years of data (1991-2020) to calibrate our model.

The average 3-Month Treasury Bill is 2.5%, which we round to 3%. We derive the mean and volatility of public equity from the S&P 500 index log returns at monthly frequency. The average log-return is 8.1%, which we round to 8%, and the volatility is 14.7%, which we round up to 15%. To sum up, we have $r = 0.03$, $\mu = 0.08$, and $\sigma = 0.15$. In addition, we use standard values for the investor risk aversion and discount factor: $\gamma = 4$ and $\delta = 0.5$. The discount for PE-funds secondary market sales is set to 13.8%, which is the average reported in Nadault et al. (2019).

One of the contributions of this paper is the detailed calibration of the private equity return dynamics. We use the Preqin dataset with private equity fund cash flows as of the end of year 2020. We select all private equity funds (venture capital, growth equity, leveraged buyout) focused on the US and raised between 1991 and 2015, hence thirty year of data.

We construct two cumulative distribution functions for PE-funds cash inflows: i) the empirical one, and ii) the model-implied one. The former is derived directly from the Preqin dataset. Preqin records all cash inflows into any given private equity fund, without a distinction between fee payments and capital invested.

The latter is generated by using 100,000 simulations of our model: a PE-fund cash *inflows* consist of the management fees from the time of capital commitment (τ_0) to the time of investment (τ_C), plus the capital called for investment at time τ_C.

Note that the time from τ_0 to τ_C follows an exponential distribution. Therefore, at any time t, the proportion of funds across simulations having not called yet is $e^{-\lambda_C t}$. The total fee paid by these funds is $(fee)(e^{-\lambda_C t})$; where $fee = 2\%$. The cumulative cash inflow across simulations at time t is the sum of the cumulative fees paid by each fund until its capital call, $\int_{\tau_0}^{t}(fee)(e^{-\lambda_C t})dt = \frac{fee}{\lambda_C}\left(1 - e^{-\lambda_C t}\right)$, and of the amount of capital already called $1 - e^{-\lambda_C t}$. The total inflow thus follows an exponential distribution with parameter λ_C.

We search for the λ_C that minimizes the least-square distance between this model-implied and the empirical cumulative distribution. 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 functions are very close to one another, which validates our modeling choice.

To calibrate the PE 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. Hence, we would be restricted to using a sample of funds raised before 2005.

Instead, we use the same sample as above, hence funds raised up to 2015, and match the distribution of their performance as of the end of the year 2020, valuing their unexited investments at their reported Net Asset Value (NAV), as done in practice. To measure fund performance, we also 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. Finally, the empirical cumulative distribution of PMEs is weighted by fund size in order to have a more economically meaningful performance measure (compared to equal-weight).

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. There are four free parameters in our model: the intensity of capital distributions, λ_D, the correlation between private and public equity (ρ), PE expected return (ν) and PE volatility (ϕ).

As with the calibration above, we choose the four free parameters $(\nu, \phi, \lambda_D, \rho)$ which minimize the least-square distance between the model-implied and the empirical cumulative distributions. The best fit is obtained with a combination of a relatively high ν and ϕ:

$\nu = 14.5\%$; $\phi = 24.7\%$; $\rho = 80\%$; $\lambda_D = 0.227$

We therefore assume an expected PE return that is 6.5% higher than that of public equity (using the S&P 500 index) and a volatility of PE is 63% higher than that of public equity.[17]

[17] The interpretation of volatility is not trivial. The value-weighted distribution of PMEs has a volatility of about 50% and the volatility of the GBM depends primarily (but not only) of this parameter. If we had equally weighted PMEs, this cross-sectional volatility would have much larger (about 80%). If we create portfolios of funds, the volatility decreases with the number of funds in the portfolio. With 5 funds in a portfolio we would have the same volatility as the value weighted one (40%). Most importantly, as shown below, this implied volatility of the GBM leads to a reasonable estimate of Beta, which leads to a reasonable allocation to PE. Also, the usual approach to estimate volatility consists in using a time-series of NAV-to-NAV IRRs at a quarterly frequency. These time-series are readily available; for example, on the website of Cambridge Associates. The resulting volatility is typically about half the volatility of public equity due to the well-known staleness of NAVs (see, e.g., Stafford (2020)). Using standard de-smoothing approaches (i.e. assuming an autoregressive process) brings the volatility closer to that of public equity, which is surprising.

The calibrated λ_D implies an average holding period nearly 4.5 years, which matches both anecdotal evidence and large sample evidence (Lopez de Silanes et al. 2011). The correlation of 80% also matches estimates that are used in practice.[18] Notice that these parameters produce an implied private equity β that matches empirical estimates obtained with direct approaches (e.g. Ang et al. (2018)):[19]

$\beta = \rho.\frac{\sigma}{\psi} = 0.8(\frac{0.25}{0.15}) = 1.32$

The value-weighted cumulative distribution of PMEs is shown in Panel B of Figure 1. The blue bars represent the empirical distribution and the red line the model-implied distribution. The two curve are close with the partial exception of the tails.

5 Portfolio Allocation in the Baseline Economy

With the above calibration, the investor optimally commits 21.7% of her wealth to private equity (ξ). This initial PE allocation will then vary with time, following changes in the investor's wealth. Figure 2, Panels A, B, C and D plot four functions describing the solution to the model: i) the value function solving (9), ii) the probability density function of ξ, iii) the consumption rate and iv) the stock allocation; all as functions of ξ. The dashed line displays the functions during the commitment period, i.e. before the capital has been called, and the solid line during the holding period, i.e. after the capital has been called. The optimal pledge is shown as a vertical dotted line and coincides with the maximum of the value function during the commitment period. Note that these functions differ across the two sub-periods (commitment versus holding period), but not within sub-periods.

Panel A displays the value function $(-\log(-H(.))$ as a function of ξ. Interestingly, we note that the region around the optimal pledge is fairly flat, implying that as long as the PE allocation remains between 0% and 40%, welfare does not vary much. However, if the PE allocation goes above 40%, welfare is quickly reduced. If it becomes larger than 78.7% of liquid wealth (black circle), the investor will strategically default on her pledge and permanently lose access to the private equity market. However, Panel B shows that the likelihood of reaching a 40% PE allocation during the commitment period is almost zero. Hence, default hardly ever happens. In fact, the PE allocation does not vary much and

[18] `https://www.blackrock.com/institutions/en-us/insights/charts/capital-market-assumptions`. Notice as well that the estimate of the volatility of US private equity is 32%, which is close to our estimate.

[19] Ang et al. (2018) propose a methodology to estimate the β of all private equity funds in a CAPM setting. They find a β of 1.43, which is close to the β found above. The fact we have implied parameters that match existing point estimates is very reassuring.

remains smaller than 25.2% with a 90% probability during the commitment period.

Upon capital call, the value function jumps up to the solid black line. At that point in time, the optimal PE allocation is 34.3%. The investor is therefore significantly under-committed to PE at this point in time. If she had committed more to PE, however, she would have increased the probability to reach lower values of welfare, and to default.

Ideally, an investor would therefore like to commit 21.7% of her wealth to PE, and have the possibility to top-up that amount by more than 60% at the time of capital call. This result gives an order of magnitude for the amount of co-investment opportunities an investor would like to be offered. In practice, although this practice has developed substantially over the last 15 years, the usual top-up is about 10%. The asymmetry of demand is also interesting. Because the investor is under-committed to PE, she does not value much the option to reduce her allocation to PE upon capital call. In contrast, she values a lot a top-up option. The latter is exactly what co-investments offer.

After the capital call, the investor can either keep the investment until capital distribution, or sell her stakes on the secondary market at a discount. The latter is optimal if the private equity share ξ reaches 94.5% of total wealth (black square). It is therefore an extremely rare event. We believe this is a realistic feature as secondary markets are indeed rarely used for fire sale exits.

Panel C shows that the consumption function is similar to the value function. During the commitment period, consumption is relatively insensitive to the PE allocation. At the optimal PE allocation, the consumption rate is at its highest point (4.8%). If the PE allocation goes beyond 40%, the investor reduces consumption, and at the point of default ($\xi = 78.7\%$), the consumption rate is 4.5%. In contrast, during the holding period (the solid line), consumption becomes much more sensitive to the PE allocation. If the PE allocation goes beyond 50% of wealth, consumption decreases rapidly to about 4%. Overall, consumption is higher during the holding period than during the commitment period.

The stock allocation (Panel D) is flat during the commitment period at about 56%, except when the PE allocation passes 60%. At this point, the stock allocation increases rapidly as the PE allocation increases, and reaches 80% of wealth. The investor is "gambling for resurrection", i.e., increasing her risky allocation with the hope to avoid default. The allocation to the risk free asset during the commitment period is therefore relatively high as long as the PE allocation stays below 60%. Despite the significant equity premium, 44% of the investor's wealth is invested in the bond, therefore the investor has almost surely enough wealth to cover for her PE commitment, since bonds are risk-free.

6 Liquidity frictions

In this section we quantify the impact of the different liquidity frictions described in Section 3.5 on the investor's welfare, her portfolio allocation and consumption rate. The cost associated with each friction is measured as the amount of wealth the investor is willing to pay to remove the given friction. We measure these costs by comparing the different economies described in Section 3.5.

6.1 Timing Risk

We switch off call timing risk by making the capital call time deterministic (Economy 1). Surprisingly, removing timing risk has almost no impact on the optimal commitment, consumption policy and stock allocation (Table 1 and Figure 3). In fact, the cost associated with commitment timing risk is negative, implying that the agent prefers uncertainty about the timing of capital calls to certainty. She is willing to pay 0.17% of her wealth to make the call stochastic, i.e. nearly 1% of what she commits to PE.

Two competing forces can explain this result. 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 induces uncertainty about the time at which the agent will realize this utility gain. As the present value $e^{-\delta \tau_C} U(1)$ of the utility of a dollar received at a random time τ_C is higher than the present value $e^{-\frac{\delta}{\lambda_C}} U(1)$ of the utility of a dollar received at the time $\mathbb{E}[\tau_C]$, the investor prefers that the time of her gain be stochastic rather than deterministic.[20] This first force pulls the cost of commitment-timing risk down.

On the other hand, a stochastic capital call time induces uncertainty in the ability of the agent to fund her requirements. These requirements include the capital call, but also the future consumption stream, taken out of liquid wealth. The larger the variance of the capital call time, the more likely it is that capital call occur far beyond $\mathbb{E}[\tau_C]$ and that she then needs to reduce consumption. This second force pushes the cost of commitment-timing risk up. The cost of timing risk depends on the relative effect of these two forces. A low subjective discount factor or low interest rate leads to the second force dominating.

To sum up, timing risk has a negligible effect and, with our default parameterization leads to a negative cost.

In contrast, distribution timing risk matters. In Economy 5, we switch off this risk, by making the time of distributions deterministic. Table 1 shows that that the optimal PE

[20] Applying Jensen's inequality to the convex function $t \to e^{-\delta t}$ yields $e^{-\delta \mathbb{E}[\tau]} \leq \mathbb{E}[e^{-\delta \tau}]$.

allocation significantly increases: from 21.7% to 25.9%. Figure 3, Panels B and D, show that the investor's consumption also increases during the commitment period. Other aspects of the economy are more similar to Economy 1. Default is triggered only slightly later in Economy 5 than in Economy 1. The stock allocation and the consumption rate are similar in both economies.

Distribution timing risk has a positive and significant cost. Our estimate of 1.07% of the investor's wealth is similar to what is documented in Ang et al. (2014). The investor thus strongly prefers a certain distribution time to an uncertain one. Certainty on the distribution timing is particularly valuable when the investor needs to reduce her consumption to avoid running out of liquid wealth by the time of capital distribution.

6.2 Quantity Risk

To evaluate the impact of quantity risk, we compare Economy 0 (baseline) to Economy 2, in which the committed amount can be chosen at capital call time. If the investor can choose the amount to invest when capital is called, it means that commitment risk is eliminated.

Table 1 shows that switching off quantity risk leads to a significant increase of the PE allocation: from 21.7% to 33.6%. The rest of variables are hardly affected: consumption increases only slightly (Figure 3, Panel C), and the stock-bond split is unaffected. The investor is willing to pay as much as 5.45% of her wealth to switch from Economy 0 to Economy 2, i.e. to be given the possibility to adjust her commitment at call time.

Switching off timing risk in Economy 2 leads us to Economy 3, in which the time of capital call is deterministic, and the investor can choose the amount she invests at call time. Removing timing risk actually decreases PE allocation (34% of wealth to 25%), and induces a small cost of 0.43%. Consumption and stock allocation are left unchanged.

In Economy 3, the investor faces a deterministic period of nearly three years during which she has no private equity exposure. We remove this period in Economy 4, by allowing her to invest directly in the fund. Her portfolio allocation is similar to the one in Economy 2. She increases her consumption by around 8%, in anticipation of the distributions that she will now perceive earlier.

6.3 Secondary Market Liquidity

The secondary market allows the investor to sell her private equity fund investment any time after the capital has been called, at a cost. Our default parameterization sets this cost (a.k.a. the haircut) to 13.8% of investment value. Table 2, Panel A, shows how the PE allocation changes if the haircut increases to 40%, or decreases to 0%.

With a perfectly liquid secondary market, PE allocation only slightly increases: from 21.7% to 22.8%. Similarly, an illiquid secondary market only decreases PE allocation to 20.4%. Furthermore, Panel B shows that the investor is not willing to pay much to increase the liquidity of the secondary market in a given economy. In the baseline economy, she is only willing to give up 0.2% of her wealth to decrease the secondary market haircut, from 13.8% to 0. This amount grows to 0.8% in economies where quantity risk is switched off and to 4.1% when getting rid of the commitment period. A secondary market is therefore more valuable if there is no ex ante commitment of capital (i.e. no quantity risk). However, these amounts remain small compared to the amounts the investor is willing to pay to switch economy, for a given secondary market haircut. The discrepancy justifies our focus on commitment risk, rather than on secondary market liquidity.

A perhaps surprising result is that the investor is willing to pay more to get rid off quantity risk with a liquid secondary market than with an illiquid secondary market, as shown in Panel C. The willingness to pay to get rid of quantity risk increases when the liquid secondary market is liquid: from 5.5% to 6%. This increase in cost is explained by an increase in PE allocations: from 33.6% to 38.3%. The willingness to pay to get rid of the commitment period also increases, from 13.8% to 17.2%. The two types of frictions, namely commitment risk and illiquidity of the secondary market, are therefore complement.

We also note that the effect is asymmetric. Changing the haircut from 13.8% to 0% had a larger effect than when we increase it from 13.8% to 40%. The cost of quantity risk goes down from 5.5% to 5%, and the PE allocation in Economy 2 goes from 33.6% to 30.7%. Making the secondary market perfectly liquid does not affect much our previous results on timing risk (allocations and cost).

6.4 Extension to Two Private Equity Funds

We now analyze an economy with two private equity funds. Their return dynamics are given by equation (3), with expected return ν and volatility ψ. The Brownian motions that drive them have correlation ρ_{PE}. As we focus on the effect of commitment risk, we assume that the two funds have almost perfectly correlated returns, and so any welfare changes compared to the exercise with one fund must be the result of commitment risk effects: $\rho_{PE} = 0.95$.

Our analysis of the two-fund case builds upon the one-fund case. The main difference is that the value function now admits two state variables and two fund states. At time zero, the investor will choose the optimal ξ_1 and ξ_2 such that the value function $H(\xi^1, \xi^2, S^1 = C, S^2 = C)$ be at its maximum.

Table 3 displays the PE allocation to each PE fund, and the welfare costs of the different economies. In the baseline economy, the PE allocation increases compared to the one-fund case: from 21.7% to 29.6%. Intuitively, with access to two funds, the investor would like to take advantage of liquidity diversification by spreading her allocation across funds. Doing so allows her to smooth her capital inflows and outflows. As with all risk-averse investors, smoothing wealth flows results in higher risky asset allocation.

A surprising result is that the investor is willing to pay more than in the one-fund case to get rid of quantity risk, even though she does not change her PE allocation much: the welfare cost increases from to 5.5% to 6.4%. This result too goes against the intuition of liquidity diversification.

Robinson and Sensoy (2013) show that the timing of capital calls and distributions is mostly idiosyncratic. In our model these times are fully idiosyncratic, as calls and distributions are driven by independent Poisson processes. Intuitively, access to two funds should therefore induce liquidity diversification, and decrease the welfare costs of the different liquidity frictions. Practically, with more than one fund the investor can space out investments to reduce the impact of commitment delay. Funding capital calls should also be facilitated, as the distribution of the first fund can be used to fund the capital call of the second fund.

However, this logic does not account for the funding mismatch that arises in the two-fund case. The randomness of capital call and payout times raises uncertainty on the funding needs of these funds. When the investor has two PE funds, she runs the risk that the second fund calls her pledge before the first fund returns its capital. This effect is not possible with one fund. As a result of this possible funding mismatch, a larger endogenous 'escrow' account is needed, which increase the welfare cost of commitment-quantity risk.

References

Ang, A., B. Chen, W. N. Goetzmann, and L. Phalippou (2018). Estimating private equity returns from limited partner cash flows. *Journal of Finance 73*, 1751–1783.

Ang, A., D. Papanikolaou, and M. M. Westerfield (2014). Portfolio choice with illiquid assets. *Management Science 60*(11), 2381–2617.

Banal-Estañol, A., F. Ippolito, and S. Vicente (2017). Default penalties in private equity partnerships. Working paper.

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

Cox, J. C. and C. Huang (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. *Journal of Economic Theory 49*(1), 33 – 83.

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

Dai, M., P. Li, H. Liu, and Y. Wang (2015). Portfolio Choice with Market Closure and Implications for Liquidity Premia. *Management Science 62*(2), 368–386.

Dimmock, S. G., N. Wang, and J. Yang (2019). The Endowment Model and Modern Portfolio Theory. Working paper.

Dybvig, P. H. and C. Huang (1988). Nonnegative wealth, absence of arbitrage, and feasible consumption plans. *Review of Financial Studies 1*(4), 377–401.

Fang, L., V. Ivashina, and J. Lerner (2015). The disintermediation of financial markets: Direct investing in private equity. *Journal of Financial Economics 116*(1), 160–178.

Franzoni, F., E. Nowak, and L. Phalippou (2012). Private Equity Performance and Liquidity Risk. *Journal of Finance 67*(6), 2341–2373.

Giommetti, N. and M. Sorensen (2020). Optimal Allocation to Private Equity. Working paper.

Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time* (Second ed.), Volume 24 of *Applications of Mathematics*. New York: Springer-Verlag.

Lerner, J., J. Mao, A. Schoar, and N. R. Zhang (2020). Investing Outside the Box: Evidence from Alternative Vehicles in Private Equity. *Journal of Financial Economics*. Forthcoming.

Lerner, J. and A. Schoar (2004). The illiquidity puzzle: theory and evidence from private equity. *Journal of Financial Economics 72*(1), 3–40.

Ljungqvist, A., M. Richardson, and D. Wolfenzon (2017). The Investment Behavior of Buyout Funds: Theory and Evidence. *Working paper*, NYU.

Longstaff, F. A. (2001, June). Optimal Portfolio Choice and the Valuation of Illiquid Securities. *Review of Financial Studies 14*(2), 407–431.

Longstaff, F. A. (2009). Portfolio Claustrophobia: Asset Pricing in Markets with Illiquid Assets. *American Economic Review 99*(4), 1119–1144.

Metrick, A. and A. Yasuda (2010). The Economics of Private Equity Funds. *Review of Financial Studies 23*(6), 2303–2341.

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

Officer, M. S., O. Ozbas, and B. A. Sensoy (2010). Club deals in leveraged buyouts. *Journal of Financial Economics 98*(2), 214–240.

Pastor, L. S. and R. F. Stambaugh (2003). Liquidity Risk and Expected Stock Returns. *Journal of Political Economy 111*(3), 642–685.

Robinson, D. T. and B. A. Sensoy (2016). Cyclicality, Performance Measurement, and Cash Flow Liquidity in Private Equity. *Journal of Financial Economics 122*(3), 521–543.

Sorensen, M., N. Wang, and J. Yang (2014). Valuing Private Equity. *Review of Financial Studies 27*(7), 1977–2021.

Stafford, E. (2020). Replicating Private Equity with Value Investing, Homemade Leverage, and Hold-to-Maturity Accounting. Working paper.

A Model solutions

A.1 Baseline model solution - 1 fund

In the case where preferences are given by power utility, the investor solves problem (6). The function F solves the HJB equation (7) and the value function F can be written as (9). We give below the reduced HJB equation solved by H.

Proposition 2 (Baseline, one fund) *The investor's value function can be written as in (9), where $H(\xi, S)$ exists and is finite, continuous, and concave for $\xi \in [0, 1)$. Whenever the investor can pledge, she selects $\xi^* \equiv \arg\max_\xi H(\xi, S = C)$, which exists. Between private equity pledges and capital calls, $H(\xi, S = C)$ is characterized by*

$$
\begin{aligned}
0 = \max_{c,\theta} \Bigg[& \frac{c^{1-\gamma}}{1-\gamma} - \beta H + A_0(c, \theta, S = C)H + A_1(\xi, c, \theta, S = C)H_\xi \\
& + A_2(\xi, c, \theta, S = C)H_{\xi\xi} + \lambda_C \left(H^{+C} - H\right)\Bigg] \qquad (11)
\end{aligned}
$$

where

$$
\begin{aligned}
A_0(c, \theta, S = C) = & \ (1-\gamma)(r + \theta(\mu - r) - c) + \frac{\gamma}{2}(\gamma - 1)\sigma^2\theta^2 \\
A_1(\xi, c, \theta, S = C) = & \ \left[-(r + \theta(\mu - r) - c) + \sigma^2\theta^2\gamma\right]\xi \\
A_2(\xi, c, \theta, S = C) = & \ \frac{1}{2}\sigma^2\theta^2\xi^2 \\
H^{+C}(\xi, S = C) = & \ H(\xi, S = D).
\end{aligned}
$$

Between capital calls and distributions, $H(\xi, S = D)$ is characterized by

$$
\begin{aligned}
0 = \max_{c,\theta} \Bigg[& \frac{[c(1-\xi)]^{1-\gamma}}{1-\gamma} - \beta H + A_0(\xi, c, \theta, S = D)H + A_1(\xi, c, \theta, S = D)H_\xi \\
& + A_2(\xi, c, \theta, S = D)H_{\xi\xi} + \lambda_D(\max_\xi H^{+D} - H)\Bigg] \qquad (12)
\end{aligned}
$$

where

$$
\begin{aligned}
A_0(\xi, c, \theta, S = D) = & \ (1-\xi)(1-\gamma)(r + \theta(\mu - r) - c) + \xi(1-\gamma)\nu \\
& + \frac{\gamma}{2}(\gamma - 1)\left(\xi^2\psi^2 + \sigma^2\theta^2(1-\xi)^2 + 2\xi(1-\xi)\rho_L\psi\sigma\theta\right) \\
A_1(\xi, c, \theta, S = D) = & -\xi(1-\xi)(r + \theta(\mu - r) - c) + \xi(1-\xi)\nu \\
& + \gamma\left(-\psi^2\xi^2(1-\xi) + \sigma^2\theta^2(1-\xi)^2\xi - \xi(1-\xi)(1-2\xi)\rho_L\psi\sigma\theta\right) \\
A_2(\xi, c, \theta, S = D) = & \ \frac{1}{2}\xi^2(1-\xi)^2\left(\psi^2 - 2\rho_L\sigma\theta\psi + \sigma^2\theta^2\right) \\
H^{+D}(\xi, S = D) = & \ H(\xi, S = C)
\end{aligned}
$$

If the investor has Epstein-Zin preferences and solves Problem (**??**), the function H^{EZ} solves analogous ODEs.

Proposition 3 (Baseline Epstein-Zin, one fund) *The investor's value function can be written as in (9). Between private equity pledges and capital calls, $H^{EZ}(\xi, S = C)$ is characterized by*

$$0 = \max_{c,\theta} \Bigg[\frac{\beta}{1-\zeta} \left(\frac{c^{1-\zeta}}{((1-\gamma)H^{EZ})^{\frac{\gamma-\zeta}{1-\gamma}}} - (1-\gamma)H^{EZ} \right) + A_0(c,\theta,S=C)H^{EZ} \\ + A_1(\xi,c,\theta,S=C)H^{EZ}_{\xi} + A_2(\xi,c,\theta,S=C)H^{EZ}_{\xi\xi} + \lambda_C \left(H^{EZ,+C} - H^{EZ}\right) \Bigg] \tag{13}$$

Between capital calls and distributions, $H(\xi, S = D)$ is characterized by

$$0 = \max_{c,\theta} \Bigg[\frac{\beta}{1-\zeta} \left(\frac{[c(1-\zeta)]^{1-\zeta}}{((1-\gamma)H^{EZ})^{\frac{\gamma-\zeta}{1-\gamma}}} - (1-\gamma)H^{EZ} \right) + A_0(\xi,c,\theta,S=D)H^{EZ} \\ + A_1(\xi,c,\theta,S=D)H^{EZ}_{\xi} + A_2(\xi,c,\theta,S=D)\frac{1}{2}H^{EZ}_{\xi\xi} + \lambda_D(\max_{\xi} H^{EZ,+D} - H^{EZ}) \Bigg] \tag{14}$$

The functions A_0, A_1 and A_2 are defined as above.

B Numerical Methods

After changing variables: $\xi = e^Z$ when $S = C$ and $\xi = \frac{1}{1+e^{-Z}}$ when $S = D$, the value function (9) can be rewritten as

$$F(W, X, S) = W^{1-\gamma} G(Z, S), \tag{15}$$

where $Z = \log(X/W)$, and G solves the following transformed reduced HJB equation:

$$0 = \max_{c,\theta,X} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + B_0(c,\theta)G + B_1(c,\theta)G_Z + B_2(c,\theta)G_{ZZ} + \lambda_S \left(G^{+S} - G\right) \right\}, \tag{16}$$

where $\lambda_S = \lambda_C$ (resp. $\lambda_S = \lambda_D$) when $S = C$ (resp. $S = D$). The functions B_0, B_1 and B_2 and G^{+S} are defined as follows:

$$\begin{cases} B_0(c,\theta) = -\beta + (1-\gamma)\left(r + (\mu - r)\theta - c\right) - \frac{1}{2}\gamma(1-\gamma)\theta^2\sigma^2 \\ B_1(c,\theta) = -\left(r + (\mu-r)\theta - c\right) + \theta^2\sigma^2(\gamma - \frac{1}{2}) + \nu - \frac{1}{2}\psi^2 + \theta\sigma\psi\rho_L(1-\gamma) \\ B_2(c,\theta) = \frac{1}{2}\theta^2\sigma^2 + \frac{1}{2}\psi^2 - \theta\sigma\psi\rho_L \\ G^{+S} = G(Z, S=D) \text{ if } S = C \text{ and } G^{+S} = G(Z, S=C) \text{ if } S = D. \end{cases} \tag{17}$$

When $S = C$, the parameters ρ_L, ψ and ν are set equal to zero.

We use the Markov Chain Approximation (MCA) of Kushner and Dupuis (2001) to evaluate equations (16), for $S = C$ and $S = D$. The method used applies an MCA to continuous-time continuous state stochastic control problems, by renormalizing finite differences forms as proper Markov chain transition probabilities. These transition probabilities arise when deriving finite difference versions of the HJB equation. In our case, the ODEs (PDEs with more than one fund) are linked: $G(Z, S = C)$ depends on $G(Z, S = D$ and vice versa, which introduces additional complexity.

We present below the main steps of the algorithm, for one private equity fund and with power utility.

We define a grid z_i of values taken by Z, ranging from -10 to 5 with intervals of $dz = \frac{1}{100}$. For each state S, we discretize the value function G on the grid of z_i: $g_i = G(z_i)$. g_i satisfies an equation of the form:

$$0 = \max_{c,\theta} \Bigg[\tilde{U}(z_i, c, \theta) + B_0(z_i, c, \theta) g_i + B_1^+(z_i, c, \theta) \frac{g_{i+1} - g_i}{h} + B_2^-(z_i, c, \theta) \frac{g_i - g_{i-1}}{h} + B_2(z_i, c, \theta) \frac{g_{i+1} + g_{i-1} - 2g_i}{h^2} \Bigg],$$

where $\tilde{U}(z_i, c, \theta)$ includes the utility term as well as the jump term, $B_1^+(z_i, c, \theta)$ (resp. $B_1^-(z_i, c, \theta)$ is the positive (resp. negative) part of $B_1(z_i, c, \theta)$.

The optimal allocation and consumption policy, as functions of f_i, are obtained from the first order condition.

We rearrange the terms to define transition probabilities and the time step of a Markov chain:

$$\begin{cases} p_u(z_i, c, \theta) = \dfrac{\frac{1}{h} B_1^+(z_i, c, \theta) + \frac{1}{h^2} B_2(z_i, c, \theta)}{-B_0(z_i, c, \theta) + \frac{1}{h} B_1^+(z_i, c, \theta) - \frac{1}{h} B_1^-(z_i, c, \theta) + \frac{2}{h^2} B_2(z_i, c, \theta)} \\[2ex] p_d(z_i, c, \theta) = \dfrac{-\frac{1}{h} B_1^-(z_i, c, \theta) + \frac{1}{h^2} B_2(z_i, c, \theta)}{-B_0(z_i, c, \theta) + \frac{1}{h} B_1^+(z_i, c, \theta) - \frac{1}{h} B_1^-(z_i, c, \theta) + \frac{2}{h^2} B_2(z_i, c, \theta)} \\[2ex] \delta(z_i, c, \theta) = \dfrac{1}{-B_0(z_i, c, \theta) + \frac{1}{h} B_1^+(z_i, c, \theta) - \frac{1}{h} B_1^-(z_i, c, \theta) + \frac{2}{h^2} B_2(z_i, c, \theta)} \end{cases} \tag{18}$$

Conditional on the allocation and consumption policy, the value function g_i is then given by:

$$g_i = p_u(z_i, c, \theta) g_{i+1} + p_d(z_i, c, \theta) g_{i-1} + \tilde{U}(z_i, c, \theta) \delta(z_i, c, \theta). \tag{19}$$

The algorithm then consists of two steps: i) the estimation of the value function g_i, for the two states, conditional on the allocation and consumption policy, and ii) the estimation of the optimal allocation and consumption policy given the value function. These two steps are iterated until convergence. Denote the value function at iteration k by g_i^k. We stop the procedure when the sum of the absolute value of all innovations is below 10^{-6}: $\sum_i |g_i^{k+1} - g_i^k < 10^{-6}$.

Evaluating the value function involves another iteration within step i), as the PDEs for the two states of the economy are linked.

Finally, when studying economies in which there is a secondary market or the private equity fund is fully liquid, an additional iteration within the above mentioned iteration is required. Denote by h the index of that iteration. The value function is updated according to:

$$g_i^{k,h} = p_u(z_i, c^{k-1}, \theta^{k-1}) g_{i+1}^{k-1,h-1} + p_d(z_i, c^{k-1}, \theta^{k-1}) g_{i-1}^{k-1,h-1} + \tilde{U}(z_i, c^{k-1}, \theta^{k-1}) \delta(z_i, c^{k-1}, \theta^{k-1}). \tag{20}$$

$g_i^{k,h}$ is then updated to account for the value of selling the fund. Once this is done, $g_i^{k,h+1}$ is calculated using (20), etc.

For two private equity funds, we use grid intervals for z_i of $\frac{1}{50}$ and a total tolerance of 10^{-3}. Because $N = 2$ require two state variables, we are solving linked-PDEs on $751 \times 751 = 564,001$ points, rather than linked-ODEs on 1501 points.

For probability distributions we use Monte Carlo methods: we use $dt = \frac{1}{100}$ ($dt = \frac{1}{50}$ for $N = 2, \infty$) and simulate wealth shocks using $dZ_t = \epsilon \sim N(0, \sqrt{dt})$. We create a single times series lasting for $1,000,000$ years, taking the evolution of wealth from the budget equations and optimal allocations and consumption from the HJB equation.

Our standard parameters are shown in Table **??**.

Table 1: Optimal Commitments

		Commitments	Welfare costs
E0	Baseline economy	21.65%	-
E1	Deterministic call time	21.44%	-0.17%
E2	Choose quantity on call	33.62%	5.45%
E3	Choose quantity on call + Deterministic call	24.79%	5.02%
E4	No commitment period	32.96%	13.82%
E5	Deterministic payout time	25.92%	1.07%

This table reports optimal commitments and welfare costs of each economy compared to the baseline economy. Each row references an economy with a different liquidity friction.

Table 2: Secondary Market Liquidity

Panel A: Commitments

		13.8% (baseline)	0% (fully liquid)	40%
E0	Baseline economy	0.217	0.228	0.204
E1	Deterministic call time	0.214	0.228	0.202
E2	Choose quantity on call	0.336	0.383	0.307
E3	Choose quantity on call + Deterministic call	0.248	0.273	0.231
E4	No commitment period	0.330	0.336	0.295
E5	Deterministic payout time	0.259	0.259	0.259

Panel B: Welfare cost of changing secondary market liquidity

		13.8% (baseline)	0% (fully liquid)	40%
E0	Baseline economy	0	0.002	-0.002
E1	Deterministic call time	0	0.002	-0.002
E2	Choose quantity on call	0	0.008	-0.007
E3	Choose quantity on call + Deterministic call	0	0.008	-0.006
E4	No commitment period	0	0.041	-0.014
E5	Deterministic payout time	0	0.000	0.000

Panel C: Welfare cost of switching from E0 to another economy, given a haircut

		13.8% (baseline)	0% (fully liquid)	40%
E0	Baseline economy	0	0	0
E1	Deterministic call time	-0.002	-0.002	-0.002
E2	Choose quantity on call	0.055	0.060	0.050
E3	Choose quantity on call + Deterministic call	0.050	0.056	0.046
E4	No commitment period	0.138	0.172	0.128
E5	Deterministic payout time	0.011	0.009	0.013

This table reports optimal commitments and welfare costs of each economy compared to the baseline economy, for different values of the secondary market haircut. Each row references an economy with a different liquidity friction.

Table 3: Two Fund Allocation and Costs

	Commitments	Welfare costs
Baseline economy	0.148	0
Choose quantity on call	0.148	0.064
No commitment period	0.138	0.142

This table reports optimal commitments and welfare costs of each economy compared to the baseline economy, when the investor has access to two private equity funds. Each row references an economy with a different liquidity friction.

Figure 1: Model validation

Panel A: Fraction of called capital

Panel B: Distribution of PME

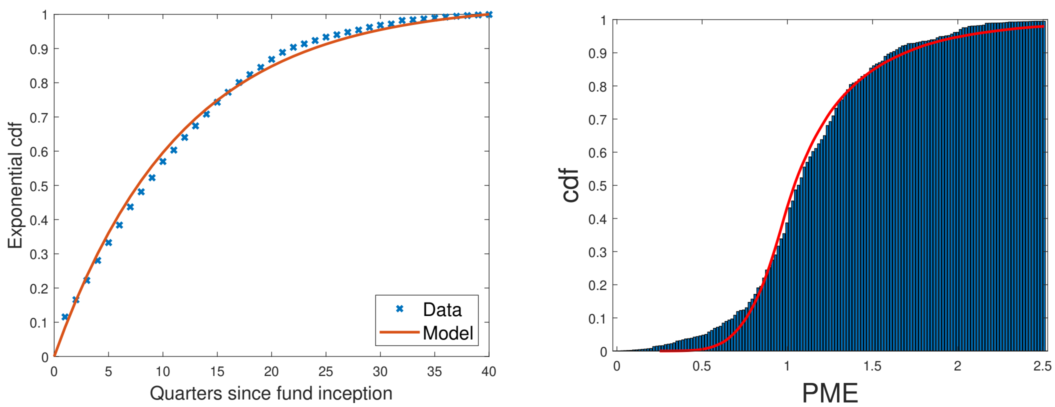

This figure illustrates the two steps of our calibration procedure for the parameters of the private equity fund dynamics. In Panel A, the blue curve represents the empirical fraction of called capital after n quarters since capital commitment, for n between 1 and 40. As the exact time of capital commitment differs across investors in a fund and is unknown, the first investment is taken as a proxy for the capital commitment time. The red curve represents the model-implied fraction of capital calls for the our calibrated λ_C of 0.344. The calibration of the intensity parameter is described in Section 4. Panel B displays the distribution of PMEs in our data sample. Panel C applies a log transformation and compares the empirical distribution of log(1+PME) to a normal distribution. The parameters ν, ϕ and λ_D are calibrated in order for the model-implied mean and standard deviation of the log(1+PME) to match their empirical counterparts.

Figure 2: Optimal Allocation and Policies in the Baseline Economy

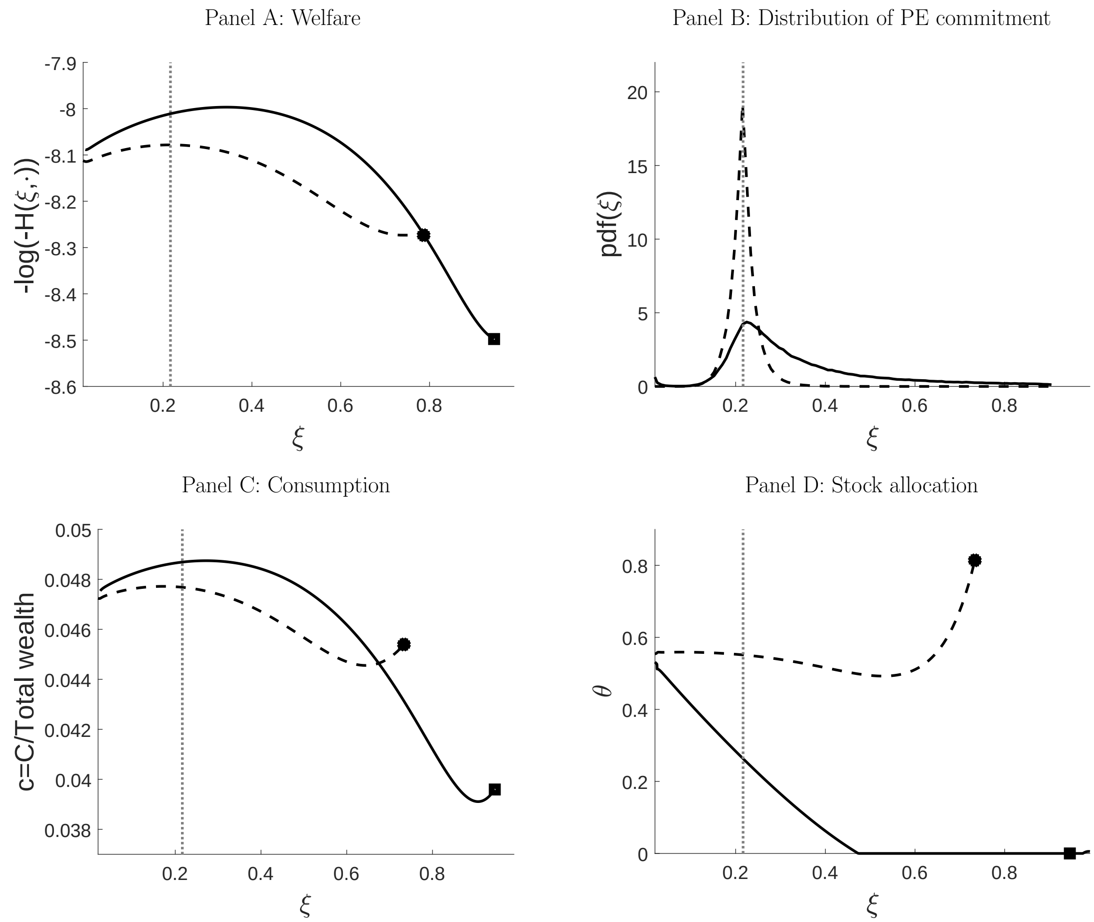

Panel A represents the value function of the investor during the commitment period (dashed line) and the holding period (plain line). Default is represented as a circle, sale on the secondary market as a square. Panel B displays the distribution of the PE allocation during the commitment period (dashed line) and holding period (plain line). Panel C displays the optimal consumption of the investor given her PE allocation. Panel D displays the optimal stock allocation.

Figure 3: Consumption and liquid asset allocation in Economies E0 to E7

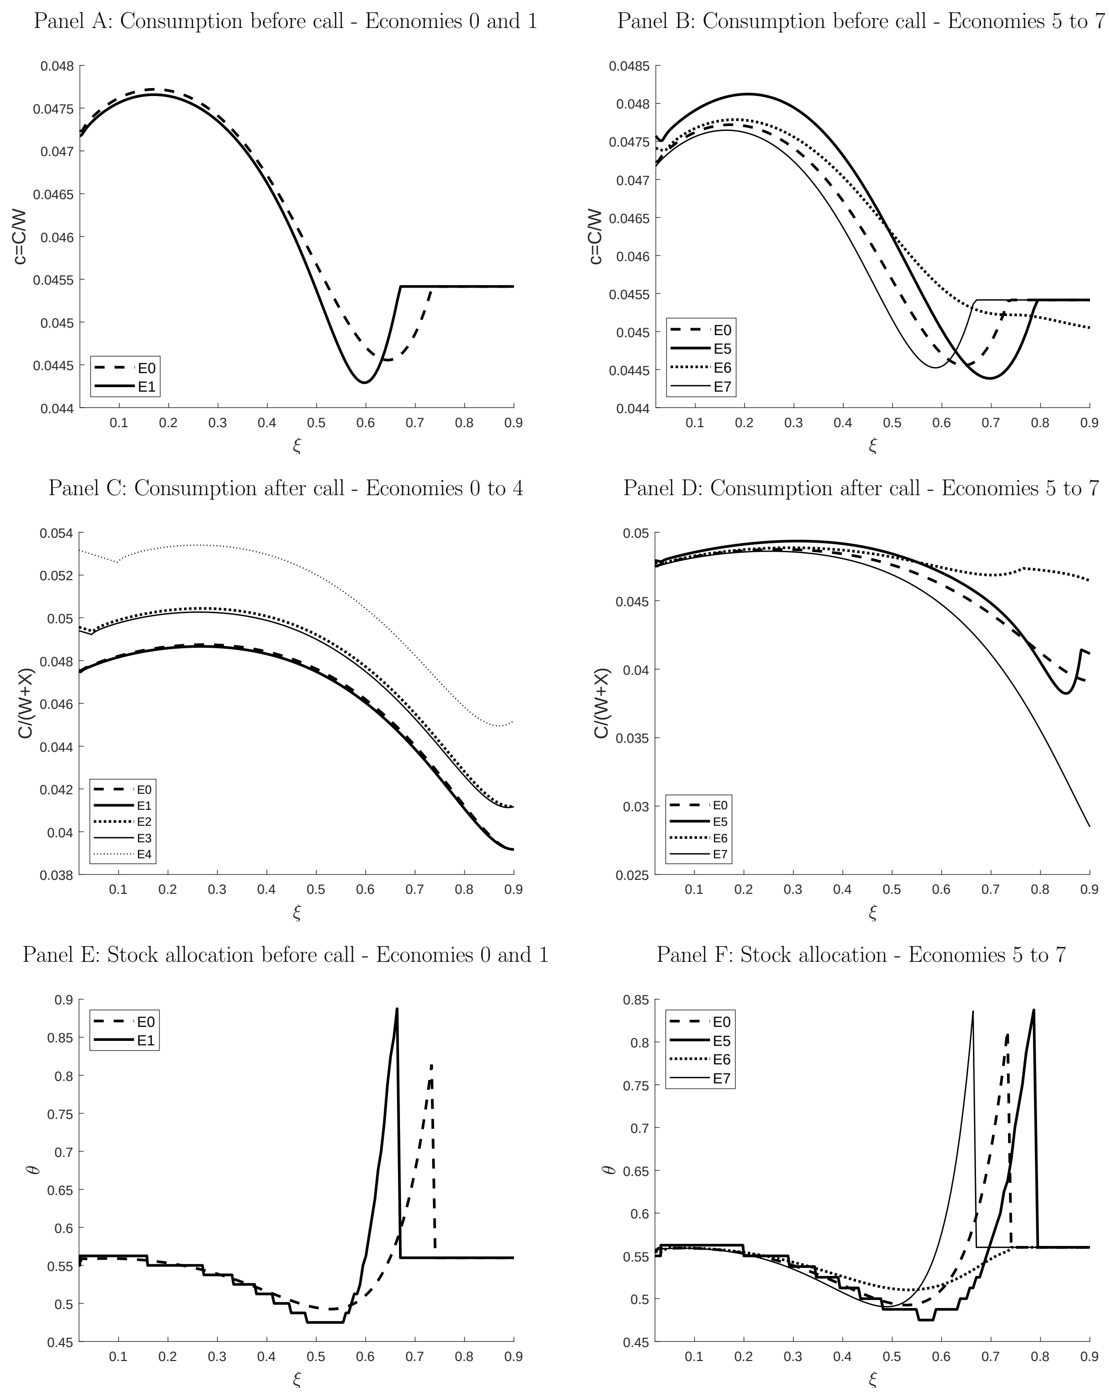

This figure represents the optimal consumption rate and liquid asset allocation in the different economies, as functions of the PE allocation. Consumption rates are given before and after capital call (Panels A to D). In Economies 2 to 4, the amount of capital invested in private equity is chosen at capital call. Stock allocation is only given before capital call (Panels E and F). After capital call, all allocations overlap, see Panel D of Figure 2.
Economies are summarized below:

Baseline economy E0	Model described in Section 3	All risks on
Economy E1	Deterministic call time	Timing commitment risk off
Economy E2	Choose quantity when called	Quantity commitment 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	Timing payout risk off
Economy E6	Fully liquid secondary market	Commitment risk on
Economy E7	Secondary market haircut at 40%	Commitment risk on