

# A Dynamic Model of Private Asset Allocation\*

Hui Chen<sup>†</sup>      Giovanni Gambarotta<sup>‡</sup>      Simon Scheidegger<sup>§</sup>

Oliver Surbek<sup>¶</sup>      Yu Xu<sup>||</sup>

December 19, 2025

## Abstract

Using cutting-edge machine learning methods, we build a state-of-the-art dynamic model of private asset allocation with five key features: (1) finite horizon; (2) illiquidity of private investments; (3) time-varying market conditions; (4) serial correlation in (observed or actual) private asset returns; and (5) regulation-induced portfolio constraints. The model identifies three distinct lifecycle stages for a private asset investor. Both the marginal value of liquidity and the relative value of illiquid assets can be highly sensitive to the state of the portfolio and market conditions. The model makes it clear that return smoothing has no impact on the long-run optimal exposure to private assets, while the illiquidity of private investments makes it difficult to exploit any variation in the actual expected returns. Finally, our model offers regulators a tool for precisely quantifying the trade-offs when setting risk-based capital charges.

JEL classification: C63, G11, G23.

Key words: alternative assets, business cycle, liquidity, machine learning, portfolio choice, private equity, return smoothing, risk-based capital.

---

\*We thank research support from the MIT-Liberty Mutual Insurance Collaboration on this project. We also thank seminar and conference participants at Cornell University, Laval University, Syracuse University, Summer School and Conference on Machine Learning for Economics and Finance, University of Delaware, and Wolfe Research's 9th Annual Global Quantitative & Macro Investment Conference for valuable comments and suggestions.

<sup>†</sup>MIT Sloan and NBER, huichen@mit.edu

<sup>‡</sup>Liberty Mutual Investments, giovanni.gambarotta@lmi.com

<sup>§</sup>Department of Economics, HEC Lausanne, simon.scheidegger@unil.ch

<sup>¶</sup>Department of Economics, HEC Lausanne, oliver.surbek@unil.ch

<sup>||</sup>Lerner College of Business and Economics, University of Delaware, yuxu@udel.edu

# 1 Introduction

Recent estimates put the size of global private markets' assets under management at \$13.1 trillion as of June 30, 2023 ([McKinsey & Company, 2024](#)). A multitude of factors make private asset allocation decisions particularly challenging. First, the illiquid nature of private assets means that it takes considerable time and costs to establish, rebalance, and exit from a position, making careful long-term planning an integral part of private asset management. Second, this illiquidity can lead to mark-to-market induced serial correlation in observed private asset returns, which further complicates dynamic allocation decisions. Third, private market investments require investors to first commit capital and subsequently manage the liquidity needs arising from future capital calls. Fourth, some institutional investors such as insurance companies face additional regulatory constraints on their allocation choices. Fifth, investors must take all of these factors into account as market and business cycle conditions constantly change.

We develop a state-of-the-art dynamic model of private asset allocation. It captures (1) the life-cycle implications of private equity investments, including the ramp-up phase, the maintenance phase, and the wind-down phase; (2) the illiquidity inherent in private asset investing, which is highly sensitive to the state of the portfolio and market conditions; (3) time-varying market and liquidity conditions; (4) serial correlation in (observed or actual) private asset returns; and (5) regulation-induced portfolio constraints. The interplay among these factors introduces substantial nonlinearities and expands the state space, rendering the model numerically intractable with traditional solution methods. To address this challenge, we use advanced machine learning techniques—Deep Kernel Gaussian Processes ([Wilson, Hu, Salakhutdinov, and Xing, 2016](#))—to solve the underlying dynamic programming problem. The solution quantifies the optimal private asset allocation policies over the life cycle of a fund, and accounts for the multitude of factors described above. To the best of our knowledge, (i) our model is the most comprehensive private asset allocation model to date in terms of the range of factors considered, and (ii) we are the first to use Deep Kernel Gaussian Processes in the context of solving dynamic models in economics and finance.

We solve a dynamic portfolio optimization problem for a risk-averse limited partner (LP) investor with a fixed horizon. This setup differs from the literature, which typically assume investors have infinite horizon.<sup>1</sup> While the infinite horizon setup is a good description for the problem that long-horizon investors such as endowments or pension funds would face, the

---

<sup>1</sup>See, e.g., [Ang, Papanikolaou, and Westerfield \(2014\)](#); [Gourier, Phalippou, and Westerfield \(2024\)](#); [Giommetti and Sorensen \(2024\)](#).

finite-horizon setting allows us to examine life-cycle implications for investors focusing on wealth accumulation instead of consuming out of their portfolios. The LP chooses among stocks, bonds, and a private asset fund, that we will refer to as private equity (PE), subject to various constraints. To invest in the PE fund, the LP must first make capital commitments before capital is called and distributed from the PE fund. The illiquid nature of PE assets makes liquidity management central: the LP must finance capital calls out of its liquid wealth (i.e., stocks and bonds). Default occurs if the LP does not have sufficient liquid wealth to meet a capital call, in which case the LP's PE stake is sold at a discount on the secondary market and the LP loses future access to the PE fund. In addition to such liquidity-driven defaults, we also allow for strategic defaults on capital calls.

Our model further incorporates three realistic features that make PE investing challenging. First, a unique feature of alternative investments is the positive autocorrelation in observed returns ([Geltner, 1991](#); [Getmansky, Lo, and Makarov, 2004](#)). While it is possible that the expected returns for private assets are persistent, a primary reason is return smoothing. As [Getmansky, Lo, and Makarov \(2004\)](#) argue, sparsely-traded illiquid assets can lead to reported PE fund returns that appear smoother than their true economic counterparts, owing to mark-to-market practices. To capture these effects, we embed a return-smoothing model into our PE return structure, allowing observed PE returns to depend on past lags of true economic returns. Consequently, the LP must make allocation decisions in the presence of both persistence in the true expected returns and mark-to-market induced smoothing, as is the case in reality.

Second, key determinants of PE investing depend on time-varying market and liquidity conditions. For example, PE returns, capital calls, and distributions can all vary over the business cycle (see, e.g., [Neuberger Berman 2022](#)). We capture these time-varying conditions by modeling a macroeconomic state that evolves according to a Markov chain. All our model parameters, including those governing PE returns, call rates, distributions, and secondary market discounts, can depend on this macroeconomic state. Consequently, our framework provides guidance for optimal PE allocations over different phases of the business cycle.

Third, many PE investors are institutional investors that must comply with regulatory constraints on their asset holdings. For example, U.S. insurers' portfolio allocations are governed by Risk-Based Capital (RBC) standards overseen by the National Association of Insurance Commissioners (NAIC), while European insurers are subject to analogous RBC requirements under Solvency II (see [Eling and Holzmuller 2008](#) for an overview of RBC standards). Violations of these requirements can be extremely costly because they trigger regulatory interventions: in severe cases, an insurer may be placed under regulatory control,

which can lead to rehabilitation or even liquidation. We capture these regulatory constraints by imposing a risk budget on the LP’s portfolio, applying separate risk charges to stocks, bonds, and PE. As long as the LP’s total risk charge remains below a specified threshold, no costs arise; however, costs escalate once that threshold is breached. This modeling approach aligns with risk-charge frameworks commonly used by rating agencies to assess insurers’ RBC adequacy (see, e.g., [S&P Global 2023](#)).

The comprehensive set of factors considered in our model makes it challenging to solve the LP’s dynamic portfolio optimization problem for two main reasons. First, the interactions among these factors introduce substantial nonlinearities that must be resolved to obtain accurate allocation policies. Second, they enlarge the state space, making traditional methods infeasible due to their vulnerability to the curse of dimensionality.

To overcome these challenges, we employ Deep Kernel Gaussian Processes (DKGPs) to track the LP’s value and policy functions within our dynamic optimization framework. As proposed by [Wilson, Hu, Salakhutdinov, and Xing \(2016\)](#), DKGPs combine the strengths of Gaussian Processes (GPs) and Neural Networks (NNs). While both GPs and NNs can address high-dimensional challenges, GPs are generally more cost-effective to train since they involve fewer hyperparameters. This is especially relevant in our setting where each sample point requires solving a constrained optimization problem with multi-dimensional controls and adaptive quadrature to accurately compute expectations. However, standard GPs often lack the flexibility of NNs in capturing complex nonlinear behavior. DKGPs resolve this issue by embedding NNs into GPs, thereby harnessing the ability to model highly nonlinear dynamics while reducing the need for expensive sample points. While we use DKGPs to solve our model, this methodology may be of independent interest, as many models in economics and finance face similar challenges.

Our calibration assumes a 10-year investment horizon for the LP, and we use PE data from Liberty Mutual Investments to calibrate our model. The calibrated model yields several quantitative insights into the optimal management of a portfolio of public and private assets.

First, the model identifies three distinct lifecycle stages for a private equity investor: the initial ramp-up phase, the maintenance phase, and the final wind-down phase. Starting with no PE exposure, the LP enters the initial stage by aggressively committing capital, with the average new commitment approaching 100% of total wealth in year one. With almost no exposure to PE to start with, the LP uses substantial stock allocations to reach target risk levels. The PE exposure of the portfolio rises rapidly. By the end of Year 4, the portfolio reaches the maintenance phase, with the average PE exposure rising from 0 to 60%, while the average annual rate of new commitment drops to less than 5%. During the

same time, the share of public stock investment in total assets falls from close to 60% to below 10%. The maintenance phase lasts until Year 8, after which the new commitment rate drops gradually towards zero. These findings demonstrate that optimal allocations evolve significantly, requiring a shift from aggressive ramping to steady-state preservation to manage the complex timing lags and illiquidity inherent in private markets.

Although similar outcomes are observed in practice, the complexity of the problem has traditionally led practitioners to rely on heuristics. For instance, standard approaches often use static mean-variance analysis combined with heuristic adjustments for capital commitment lags (see, e.g., [Takahashi and Alexander 2002](#)). Our contribution is to rigorously quantify the optimal policies in a realistic setting, moving beyond rules of thumb to help private investors achieve improved outcomes. Indeed, we demonstrate that the optimal portfolio allocations obtained under our fully dynamic framework can differ dramatically from those derived using heuristic methods.

Second, our model quantifies the value of liquidity to private investors over the lifecycle of the fund and across market and liquidity conditions. We quantify two notions of the value of liquidity. The first is the marginal value of liquidity ([Bolton, Chen, and Wang, 2011, 2013](#)), which gives the value of an additional dollar of liquid wealth to the private investors. The second is the relative price of an additional unit of illiquid private assets to an additional unit of liquid wealth ([Ang, Papanikolaou, and Westerfield, 2014](#)). When the investor is facing tight liquidity constraint (when the ratio of liquid assets to total wealth is low), an additional dollar of cash can be highly valuable in staving off default. This is especially true when the amount of uncalled commitments is high, and when the market conditions are good. Similarly, the relative price of private assets, which can either be at a premium (above 1) or a discount (below 1), is a highly nonlinear function of the share of liquid wealth and uncalled commitments.

Third, the optimal allocation accounts for business cycle conditions as follows. To first order, the optimal allocation maintains the intended transition profile for the PE portfolio, as if short-term business cycle fluctuations did not exist. This is because adjusting PE allocations can be prohibitively costly given their illiquidity, commitment lags, and long planning horizons. Instead, the LP modulates overall risk exposure over the business cycle by adjusting its public stock holdings, lowering exposure during recessions when necessary. Since transaction costs for public stocks are negligible in comparison, this approach offers a cost-effective means of managing risk over the business cycle.

Failing to account for business cycles is extremely costly. To quantify these costs, we compare the outcomes for a naive LP that ignores business cycle variation against outcomes

under the optimal policy. The naive LP adopts overly aggressive PE allocations, leading to a dramatic increase in default risk: the cumulative default frequency is 13.6% for the naive approach, compared to only 0.1% under the optimal policy. Ex-ante, this suboptimal strategy translates to a 9.3% loss in initial wealth in certainty-equivalent terms.

Fourth, we examine whether it is necessary to “unsmooth” returns before making allocation decisions. Prior research has argued that mark-to-market accounting induces serial correlation, making observed alternative asset returns appear smoother than their underlying true economic performance ([Getmansky, Lo, and Makarov, 2004](#)). This smoothing can cause private assets’ true risk to be underestimated and, consequently, lead to suboptimal portfolio allocations. As a result, some have advocated unsMOOTHING returns prior to drawing portfolio inferences.

Our model revisits this issue from the perspective of a long-term PE investor and demonstrates that the need to unsmooth returns may be a moot point for such an investor. The reasons are as follows. First, long-term investors care primarily about the characteristics of long-horizon returns, which are less affected by short-term mark-to-market fluctuations. Second, even if serial correlation is an inherent feature of true PE returns, its benefits may be negated once realistic implementation lags and adjustment costs are considered. Indeed, in our baseline calibration—where observed PE returns have a quarterly autocorrelation of 0.2—we observe little difference in outcomes between a LP that optimizes its portfolio directly using the raw returns and one that first unsmooths the return series to remove autocorrelation while preserving the long-run moments.

Fifth, our model can be used to evaluate the impact of varying risk charges on the LP’s portfolio outcomes by quantifying the associated risk–return trade-off. In our baseline calibration, we impose a 50% risk charge on both public and private equities, reflecting representative industry figures (see, e.g., [S&P Global 2023](#), Table 14). Because these charges differ across asset classes, we also consider a scenario in which the risk charge is increased to 100%, near the upper end of observed ranges. Under this higher risk charge scenario, the calibrated model indicates that long-run realized returns decline from 8.4% to 7.1%, while long-run realized volatility decreases from 2.78% to 2.35%.<sup>2</sup> These findings underscore how our framework can help both institutional investors and regulators assess the cost and benefits of alternative risk-charge levels.

---

<sup>2</sup>Note that the seemingly low standard deviation is because we are reporting the standard deviation of annualized long-horizon returns over the LP’s ten year investment horizon. For example, if annual returns  $r_{t+k}$  are iid with a volatility of  $\sigma_1$ , then the standard deviation of annualized returns over a horizon of  $H$  years is  $\sigma\left(\frac{1}{H}(r_{t+1} + r_{t+2} + \dots + r_{t+H})\right) = \sigma_1/H$ .

**Related literature.** We contribute to two strands of literature. First, we advance the research on optimal private asset allocation (see, e.g., [Korteweg and Westerfield \(2022\)](#) for a recent survey). Early industry approaches to estimating exposures and cash flows include [Takahashi and Alexander \(2002\)](#). More recent work has focused on dynamic models with stochastic shocks, and highlight two key features of private asset investing: the inherent illiquidity of these assets ([Ang, Papanikolaou, and Westerfield, 2014](#); [Sorensen, Wang, and Yang, 2014](#); [Dimmock, Wang, and Yang, 2024](#)) and the delays between capital commitments, calls, and distributions ([Giommetti and Sorensen, 2024](#); [Gourier, Phalippou, and Westerfield, 2024](#)). We build on these prior works by incorporating additional factors that pose significant challenges for PE investors, including time-varying market and liquidity conditions, serial correlation in observed PE returns, and regulation-induced portfolio constraints. To the best of our knowledge, our model is the most comprehensive stochastic model of private asset allocation to date.

Second, we contribute to a growing literature that uses machine learning methods to tackle portfolio choice problems featuring higher dimensions, realistic trading frictions, and complex constraints. Most of this work has focused on liquid asset markets such as stocks and government bonds (see, e.g., [Gaegau, Scheidegger, and Trojani 2023](#); [Duarte, Duarte, and Silva 2024](#)), leaving alternative asset classes like private equity understudied in comparison—despite their rising importance in institutional portfolios. Our contribution is to extend machine learning techniques to the alternative asset allocation setting, enabling us to characterize optimal private asset allocation under a range of realistic features thereby offering practitioners a means of achieving improved outcomes. To the best of our knowledge, we are the first to deploy Deep Kernel Gaussian Processes for solving dynamic models in economics and finance. This methodological advance is of independent interest, as similar challenges frequently arise in various settings in economics and finance. For example, [Duarte, Fonseca, Goodman, and Parker \(2022\)](#) apply neural networks to a portfolio choice problem in the household context that incorporates multiple realistic features. While their methodology characterizes optimal portfolio choice under full commitment, our approach instead characterizes the optimal time-consistent policy.

## 2 Model

Time is discrete and runs from  $t = 0$  to the terminal date  $t = T$ . The limited partner (LP) investor maximizes utility over terminal wealth  $W_T$ ,

$$u(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma}, \quad (1)$$

where  $\gamma$  is the coefficient of relative risk aversion. At each time period  $t$ , the LP can invest in a risk-free bond, a public stock index, and a private asset. For simplicity, we will refer to this private asset as private equity (PE). The timing of the decisions is illustrated in Figure 1 and is described below.

**Time-varying macroeconomic conditions.** The LP's investment opportunities set varies over the business cycle. We capture time-varying business cycle conditions through a Markov process  $s_t \in \{1, 2\}$  that takes two values corresponding to recessions ( $s_t = 1$ ) and booms ( $s_t = 2$ ). We denote by  $p_{ss'} = \text{prob}(s_{t+1} = s' | s_t = s)$  the probability of transitioning from state  $s$  to state  $s'$ .

**Liquidity constraint.** The LP enters period  $t$  with liquid wealth  $W_t \geq 0$ , uncalled PE commitments  $K_t \geq 0$ , and illiquid wealth  $P_t \geq 0$ . The latter corresponds to the net asset value (NAV) of the LP's previously called PE investments. The LP then makes three choices: (1) new PE commitments  $N_t \geq 0$ , and its allocations to (2) stocks  $S_t$  and (3) bonds  $B_t$ . These choices are subject to the following liquidity constraint:

$$S_t + B_t + \gamma_N(W_t + P_t) \left( \frac{N_t}{W_t + P_t} - \bar{n} \right)^2 + \gamma_S(W_t + P_t) \left( \frac{S_t}{W_t + P_t} - \bar{s} \right)^2 = W_t \quad (2)$$

where  $\gamma_N(W_t + P_t) \left( \frac{N_t}{W_t + P_t} - \bar{n} \right)^2$  and  $\gamma_S(W_t + P_t) \left( \frac{S_t}{W_t + P_t} - \bar{s} \right)^2$  are adjustment costs on new PE commitments and stocks, respectively.<sup>3</sup> That is, stock and bond allocations and adjustment costs are paid out of liquid wealth.

Additionally, we show in [Appendix B](#) that solvency concerns rule out shorting so that

$$B_t \geq 0, \quad S_t \geq 0. \quad (3)$$

---

<sup>3</sup>In [Appendix B](#), we show that the problem is homogeneous in total wealth  $W_t + P_t$ . The forms of these adjustment costs conveniently preserves this homogeneity property.

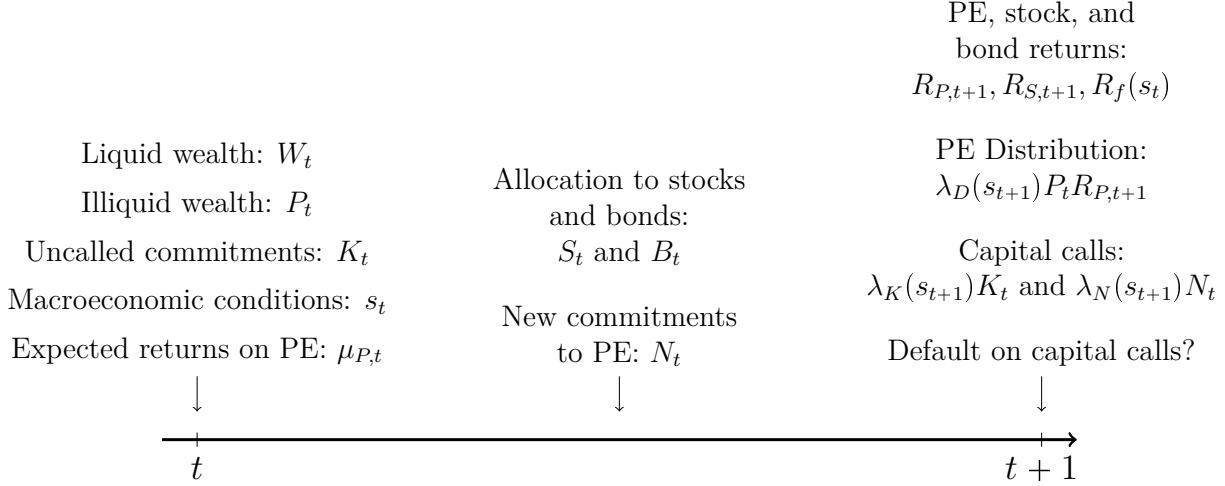


Figure 1: **Timing of the Model**

**Risk budget.** The LP's overall portfolio is subject to internal risk controls which we model through risk weights. Specifically, the risk weight of the LP's overall portfolio is defined as

$$\theta_t = \frac{\theta_B B_t + \theta_S \left[ S_t + \gamma_S (W_t + P_t) \left( \frac{S_t}{W_t + P_t} \right)^2 \right] + \theta_P P_t}{B_t + S_t + \gamma_S (W_t + P_t) \left( \frac{S_t}{W_t + P_t} \right)^2 + P_t} \quad (4)$$

The interpretation of the risk weight (4) is as follows. The LP's combined portfolio consists of  $B_t$  invested in bonds,  $S_t + \gamma_S (W_t + P_t) \left( \frac{S_t}{W_t + P_t} \right)^2$  in stocks (including stock adjustment costs), and  $P_t$  in PE. Bonds carry a risk weight of  $\theta_B$ , while stocks and PE carry risk weights of  $\theta_S$  and  $\theta_P$ , respectively. The risk weight (4) is the value-weighted average over the risk weights of the three classes of investments in the LP's portfolio. We assume the risk weights are ordered as follows:

$$0 = \theta_B < \theta_S \leq \theta_P. \quad (5)$$

That is, risk-free bonds carry a zero risk weight as they are risk free while public stocks and PE carry positive risk weights. The risk weight for PE is no lower than that of public stocks.

The LP incurs a cost if it violates its risk budget. Specifically, the cost of violating its risk budget equals  $(W_t + P_t) \times \Gamma(\theta_t)$  where the proportional cost equals

$$\Gamma(\theta_t) = \kappa (\theta_t - \bar{\theta})^2 \mathbf{1}_{\{\theta_t > \bar{\theta}\}}. \quad (6)$$

That is, no costs are assessed if the LP's risk weight (4) falls below the threshold  $\bar{\theta}$ ; quadratic costs are assessed above the threshold  $\bar{\theta}$ . Without loss of generality, we can take the threshold

to be  $\bar{\theta} = 1$ .<sup>4</sup> The risk budget cost is paid during period  $t + 1$  in a manner we describe below.

Our modeling of the risk budget through equations (4) and (6) is in line with regulatory constraints faced by institutional investors. For example, U.S. insurers' exposure to risky assets are subject to Risk-Based Capital (RBC) standards developed by the National Association of Insurance Commissioners (NAIC); similarly, the European Solvency II directive imposes RBC requirements on European insurers (see [Eling and Holzmuller 2008](#) for an overview of Risk-Based Capital standards). Such RBC requirements are also reflected in rating agencies' assessments of insurers' RBC adequacy (see, e.g., [S&P Global 2023](#)). The risk-weights  $\theta_B$ ,  $\theta_S$ , and  $\theta_P$  in our setting capture risk charges used to calculate RBC adequacy. The cost parameter  $\kappa$  can be interpreted as the shadow cost of the contribution of an insurer's PE investments to its overall RBC adequacy.

**Laws of motion.** New PE commitments affect future uncalled commitments through the law of motion

$$K_{t+1} = K_t + N_t - \underbrace{[\lambda_N(s_{t+1})N_t + \lambda_K(s_{t+1})K_t]}_{\text{capital calls}}. \quad (7)$$

New commitments  $N_t$  are added to the outstanding pool of uncalled commitments  $K_t$  before capital calls are subsequently made. A fraction  $\lambda_N(s_{t+1})$  and  $\lambda_K(s_{t+1})$  of new and previously uncalled commitments are called at the start of the next period  $t + 1$ , respectively. Both fractions depend on next period's business cycle conditions through  $s_{t+1}$ . The amount of outstanding PE commitments at the start of period  $t + 1$ ,  $K_{t+1}$ , is then given by equation (7).

The next period's illiquid wealth is determined through the law of motion

$$P_{t+1} = P_t R_{P,t+1} - \lambda_D(s_{t+1}) P_t R_{P,t+1} + \lambda_K(s_{t+1}) K_t + \lambda_N(s_{t+1}) N_t. \quad (8)$$

The NAV of the LP's PE investments at the start of period  $t$ ,  $P_t$ , earns risky gross returns  $R_{P,t+1}$  so that  $P_t R_{P,t+1}$  is the updated NAV. Two adjustments are subsequently made. First, a fraction  $\lambda_D(s_{t+1})$  of the updated NAV is distributed to the LP by its PE funds; these distributions depend on business cycle conditions. Second, the NAV increases by the LP's new PE contributions which equals capital calls  $\lambda_N(s_{t+1})N_t + \lambda_K(s_{t+1})K_t$ .

The LP's liquid wealth evolves as follows:

$$\begin{aligned} W_{t+1} = & \lambda_D(s_{t+1}) R_{P,t+1} P_t + R_{S,t+1} S_t + R_f(s_t) B_t \\ & - \lambda_K(s_{t+1}) K_t - \lambda_N(s_{t+1}) N_t - (W_t + P_t) \Gamma(\theta_t). \end{aligned} \quad (9)$$

---

<sup>4</sup>We can always rescale  $\kappa$  and the risk weights  $\theta_S$  and  $\theta_P$  appropriately to effectively enforce  $\bar{\theta} = 1$ .

The next period's liquid wealth  $W_{t+1}$  consists of distributions from its PE funds totaling  $\lambda_D(s_{t+1}) R_{P,t+1} P_t$  and the proceeds from the LP's investments in public stocks  $R_{S,t+1} S_t$  and risk-free bonds  $R_f(s_t) B_t$ . Here,  $R_{S,t+1}$  and  $R_f(s_t)$  denote the gross risky return on stocks and the gross risk-free return on bonds, respectively. Additionally,  $W_{t+1}$  is decreased by capital calls totaling  $\lambda_K(s_{t+1}) N_t + \lambda_N(s_{t+1}) N_t$  and risk budget costs  $(W_t + P_t) \Gamma(\theta_t)$ .

Finally, we assume that illiquid wealth can be converted into liquid wealth one-for-one at the terminal date  $t = T$ , and that any uncalled commitments at  $t = T$  are discarded.

**Default and liquidation.** If the investor does not have sufficient liquid wealth to meet capital calls at  $t + 1$ , it is considered a “default.” This occurs if the realized return shocks at  $t + 1$  are such that next period's liquidity (9) is negative:  $W_{t+1} < 0$ . In case of default, all uncalled commitments are written off, the LP's PE holdings are liquidated at a discount  $\alpha(s_{t+1}) < 1$ , and the investor loses access to PE investments in the future. That is, the LP enters default with only liquid wealth totaling

$$W_{D,t+1} = \lambda_D(s_{t+1}) R_{P,t+1} P_t + R_{S,t+1} S_t + R_f(s_t) B_t - (W_t + P_t) \Gamma(\theta_{D,t}) + \alpha(s_{t+1}) [1 - \lambda_D(s_{t+1})] R_{P,t+1} P_t. \quad (10)$$

The differences between the post-default liquid wealth (10) and its counterpart in the absence of default (9) are as follows. First, capital calls are not paid should the LP choose to default. Second, the liquid wealth (10) includes the proceeds from liquidating the LP's PE holdings,  $\alpha(s_{t+1}) [1 - \lambda_D(s_{t+1})] R_{P,t+1} P_t$ . Third, since PE holdings are liquidated and no longer a part of the LP's portfolio at  $t + 1$ , the risk weight (4) is modified accordingly in default:

$$\theta_{D,t} = \frac{\theta_B B_t + \theta_S \left[ S_t + \gamma_S (W_t + P_t) \left( \frac{S_t}{W_t + P_t} \right)^2 \right]}{B_t + S_t + \gamma_S (W_t + P_t) \left( \frac{S_t}{W_t + P_t} \right)^2 + P_t}. \quad (11)$$

That is, compared to the risk weight (4), the risk weight in default (11) does away with the  $\theta_P P_t$  term in the numerator. The risk cost (6) is then calculated using the risk weight in default (11).<sup>5</sup>

The LP can also choose to strategically default. Strategic default occurs when the LP has sufficient liquidity  $W_{t+1}$  at  $t + 1$  to meet capital calls, but chooses not to do so. The LP's decision to strategically default is endogenously determined as part of the LP's optimization problem and is characterized in Section 2.1.

---

<sup>5</sup>This is necessary for utilities to be well-defined—see Appendix B for details.

**Asset returns.** The gross return at  $t + 1$  for the risk-free bond  $R_f(s_t)$  depends on the macroeconomic state at time  $t$ ,  $s_t$ , whose dynamics follows a first-order Markov chain.

The gross returns for private equity  $R_{P,t+1}$  and the public stock index  $R_{S,t+1}$  follow a log-normal distribution:

$$\log \begin{bmatrix} R_{P,t+1} \\ R_{S,t+1} \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} \mu_{P,t} \\ \mu_S(s_t) \end{bmatrix}, \begin{bmatrix} \sigma_P(s_t)^2 & \rho(s_t) \sigma_P(s_t) \sigma_S(s_t) \\ \rho(s_t) \sigma_P(s_t) \sigma_S(s_t) & \sigma_S(s_t)^2 \end{bmatrix} \right). \quad (12)$$

Public stocks have expected return  $\mu_S(s_t)$ , volatility  $\sigma_S(s_t)$ , and correlation  $\rho(s_t)$  with private equity returns; all of these quantities depend on macroeconomic conditions through  $s_t$ .

The expected return to investing in private equity follows the law of motion

$$\mu_{P,t} = \varrho_{P,1} \mu_{P,t-1} + \varrho_{P,2} \log R_{P,t} + \nu_P(s_t), \quad (13)$$

where  $\nu_P(s_t)$  captures business cycle variation in private equity returns. The specification (13) allows for serial correlation in private equity returns—a hallmark feature of alternative investments (see, e.g., [Getmansky, Lo, and Makarov 2004](#)). To see this, write  $\log R_{P,t} = \mu_{P,t-1} + \sigma_P(s_{t-1}) \varepsilon_{P,t}$  where  $\varepsilon_{P,t} \equiv (\log R_{P,t} - \mu_{P,t-1}) / \sigma_P(s_{t-1}) \sim \mathcal{N}(0, 1)$  is a standard normal shock. Equation (13) can then be expressed as

$$\mu_{P,t} = (\varrho_{P,1} + \varrho_{P,2}) \mu_{P,t-1} + \varrho_{P,2} \sigma_P(s_{t-1}) \varepsilon_{P,t} + \nu_P(s_t) \quad (14)$$

from which we see  $\varrho_{P,1} + \varrho_{P,2}$  is the autocorrelation coefficient for  $\mu_{P,t}$ . In [Appendix A](#), we show that equation (14) can be motivated through the [Getmansky, Lo, and Makarov \(2004\)](#) model of smoothed returns for illiquid alternative investments.

## 2.1 Recursive Formulation

The problem has six state variables: liquid wealth  $W$ , illiquid wealth  $P$ , uncalled PE commitments  $K$ , expected PE returns  $\mu_P$ , the macroeconomic state  $s$ , and time  $t$ . In the notation that follows, we drop time subscripts and use a prime to denote variables for the next period (e.g.,  $W'$  denotes  $W_{t+1}$ ).

The problem after defaulting can be expressed recursively as follows:

$$V_D(t, W, s) = \max_{S \geq 0, B \geq 0} \mathbb{E} [V_D(t + 1, W', s') | W, s] \quad (15)$$

subject to

$$\begin{aligned} W' &= R'_S S + R_f(s) B - W \Gamma(\theta), \\ \theta &= \frac{\theta_B B + \theta_S \left[ S + \gamma_S W \left( \frac{S}{W} \right)^2 \right]}{B + S + \gamma_S W \left( \frac{S}{W} \right)^2}, \\ W &= S + B + \gamma_S W \left( \frac{S}{W} \right)^2, \end{aligned}$$

with the terminal condition being

$$V_D(T, W, s) = \frac{W^{1-\gamma}}{1-\gamma}.$$

Note that the problem after defaulting (15) does not depend on  $K$ ,  $P$ , and  $\mu_P$  because the LP loses access to PE investing after defaulting.

The problem before defaulting is given by

$$\begin{aligned} V(t, W, P, K, \mu_P, s) & \quad (16) \\ = \max_{N \geq 0, S \geq 0, B \geq 0} \mathbb{E} & \left[ \max \{ V(t+1, W', P', K', \mu'_P, s'), V_D(t+1, W'_D, s') \} \mid W, P, K, \mu_P, s \right] \end{aligned}$$

subject to

$$\begin{aligned} W' &= \lambda_D(s') R'_P P + R'_S S + R_f(s) B - \lambda_K(s') K - \lambda_N(s') N - (W + P) \Gamma(\theta), \\ P' &= [1 - \lambda_D(s')] R'_P P + \lambda_K(s') K + \lambda_N(s') N, \\ K' &= [1 - \lambda_K(s')] K + [1 - \lambda_N(s')] N, \\ \mu'_P &= \varrho_{P,1} \mu_P + \varrho_{P,2} \log R'_P + \nu_P(s'), \\ W &= S + B + \gamma_N (W + P) \left( \frac{N}{W + P} - \bar{n} \right)^2 + \gamma_S (W + P) \left( \frac{S}{W + P} \right)^2, \\ \theta &= \frac{\theta_B B + \theta_S \left[ S + \gamma_S (W + P) \left( \frac{S}{W + P} \right)^2 \right] + \theta_P P}{B + S + \gamma_S (W + P) \left( \frac{S}{W + P} \right)^2 + P}, \\ W'_D &= [\lambda_D(s') + \alpha(s') (1 - \lambda_D(s'))] R'_P P + R'_S S + R_f(s) B - (W + P) \Gamma(\theta_D), \\ \theta_D &= \frac{\theta_B B + \theta_S \left[ S + \gamma_S (W + P) \left( \frac{S}{W + P} \right)^2 \right]}{B + S + \gamma_S (W + P) \left( \frac{S}{W + P} \right)^2 + P}, \end{aligned}$$

and the terminal condition

$$V(T, W, P, K, \mu_P, s) = \frac{(W + P)^{1-\gamma}}{1 - \gamma}.$$

We additionally impose the boundary condition  $V(t, W', P', K', \mu'_P, s') = -\infty$  if  $W' < 0$  since the LP has no option but to default if it does not have sufficient liquid wealth to meet capital calls. Finally, note that the max term in the objective function of problem (16) captures the LP's strategic default decision—even with sufficient liquidity to meet capital calls, the LP can choose to default if the value of defaulting next period is higher than the value of not defaulting.

**Scaling.** The value functions (15) and (16) are homogeneous in total wealth  $W + P$  and can be scaled as follows:

$$V_D(t, W, s) = \frac{[Wv_D(t, s)]^{1-\gamma}}{1 - \gamma} \quad (17)$$

$$V(t, W, P, K, \mu_P, s) = \frac{[(W + P)v(t, w, k, \mu_P, s)]^{1-\gamma}}{1 - \gamma} \quad (18)$$

where  $w \equiv W/(W + P)$  is the liquid fraction of the LP's total wealth and  $k \equiv K/(W + P)$  is the LP's uncalled commitments relative to its total wealth.

The scaled value function  $v$  represents the certainty equivalent values for terminal wealth per unit of current total wealth, while the scaled value function  $v_D$  provides the certainty equivalent values conditional on the LP being in default. Both scaled value functions are characterized by Bellman equations, as summarized in Section B. We solve these scaled value functions using the numerical techniques described in Section 3; Section C provides further details.

### 3 Solution Technique

Our goal is to compute the global solution for the scaled value functions outlined towards the end of Section 2.1. This requires us to track the value function  $v(t, w, k, \mu_P, s)$  over time; unfortunately, the large number of state variables make it challenging to do so using traditional methods. We instead use Machine Learning methods—Deep Kernel Gaussian Process dynamic programming—to make the problem computationally feasible. We outline the key steps below and refer readers to [Rasmussen and Williams \(2005\)](#) for a textbook

treatment of Gaussian Processes (GPs), and Scheidegger and Bilionis (2019) and Renner and Scheidegger (2018) for an introduction to using GPs in a dynamic programming context.

To the best of our knowledge, we are the first to use Deep Kernel GPs in the context of solving dynamic models in economics and finance.

### 3.1 Gaussian Process Regression

Let  $f : \mathbb{R}^D \rightarrow \mathbb{R}$  be a multivariate function of interest. In our context,  $f(\cdot)$  is either a value function or a policy function. We can measure  $f(\mathbf{x})$  for input  $\mathbf{x} \in \mathbb{R}^D$  by querying an information source. We allow for noise in the information source: we measure  $y = f(\mathbf{x}) + \epsilon$  where  $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$  captures measurement noise.

In our setting, the information source is computer code that solves an optimization problem at point  $x$  in the state space. We make queries at  $N$  sample points  $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$  and observe the corresponding measurements  $\mathbf{y} = \{y^{(1)}, \dots, y^{(N)}\}$ . Computational constraints limit the number  $N$  of measurements that we can make as individual function evaluations are computationally expensive. The key idea of Gaussian Process Regression (GPR) is to replace the computationally expensive function  $f(\cdot)$  with a cheap-to-evaluate surrogate learned from the training inputs  $\mathbf{X}$  and training targets  $\mathbf{y}$ .

GPR constructs the surrogate as follows. Before making any queries, we model our prior knowledge of  $f(\cdot)$  by assigning it a GP prior:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}; \boldsymbol{\theta}), k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})).$$

That is,  $f(\mathbf{x})$  is a GP if any finite collection of function values  $\{f(\mathbf{x}_1), \dots, f(\mathbf{x}_M)\}$  has a joint Gaussian distribution with mean function  $m(\mathbf{x}; \boldsymbol{\theta}) = \mathbb{E}[f(\mathbf{x})]$  and covariance kernel  $k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}) = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}; \boldsymbol{\theta}))(f(\mathbf{x}') - m(\mathbf{x}'; \boldsymbol{\theta}))]$ . The vector  $\boldsymbol{\theta}$  denotes the hyperparameters of the GP model which must be estimated.

A typical specification for the mean function is simply a constant,

$$m(\mathbf{x}; \boldsymbol{\theta}) = m_0, \tag{19}$$

while an often-used kernel is the Matern 5/2 kernel. The automatic relevance detection (ARD) Matern 5/2 kernel is defined by

$$k_{\text{matern52}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \left( 1 + \sqrt{5}r + \frac{5}{3}r^2 \right) \exp \left( -\sqrt{5}r \right), \tag{20}$$

where

$$r = \sqrt{\sum_{d=1}^D \frac{(x_d - x'_d)^2}{\ell_d^2}}$$

and  $x_d$  is the value of  $\mathbf{x}$  along dimension  $d$ . The parameter  $\sigma_f$  modulates the output amplitude of the GP and is referred to as the signal standard deviation. The parameter  $\ell_d$  is the characteristic lengthscale of along dimension  $d$ ; it captures how quickly the GP changes as the input changes along dimension  $d$ . The resulting hyperparameters under specification (19) and (20) are  $\boldsymbol{\theta} = [m_0, \sigma_f, \ell_1, \dots, \ell_D]$ .

We use the training data  $\mathbf{X}$  and  $\mathbf{y}$  to update our beliefs regarding  $f(\mathbf{x})$ . Specifically, the posterior for  $f(\mathbf{x})$  is still a GP with posterior mean and variance given by

$$\tilde{m}(\mathbf{x}; \boldsymbol{\theta}) = m(\mathbf{x}; \boldsymbol{\theta}) + \mathbf{K}(\mathbf{x}, \mathbf{X}; \boldsymbol{\theta})^T [\mathbf{K} + \sigma_n^2 \mathbf{I}_N]^{-1} (\mathbf{y} - \mathbf{m}), \quad (21)$$

$$\text{and } \text{Var}(f(\mathbf{x})|\mathbf{X}, \mathbf{y}; \boldsymbol{\theta}) = k(\mathbf{x}, \mathbf{x}; \boldsymbol{\theta}) - \mathbf{K}(\mathbf{x}, \mathbf{X}; \boldsymbol{\theta})^T [\mathbf{K} + \sigma_n^2 \mathbf{I}_N]^{-1} \mathbf{K}(\mathbf{x}, \mathbf{X}; \boldsymbol{\theta}), \quad (22)$$

respectively. Here,  $\mathbf{K}(\mathbf{x}, \mathbf{X}; \boldsymbol{\theta})$  denotes a column vector with  $i$ th entry  $k(\mathbf{x}, \mathbf{x}^{(i)}; \boldsymbol{\theta})$ ,  $\mathbf{K}$  denotes a  $N \times N$  matrix with  $ij$ th entry  $k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}; \boldsymbol{\theta})$ ,  $\mathbf{I}_N$  denotes the  $N \times N$  identity matrix, and  $\mathbf{m}$  is a column vector with  $i$ th entry  $m(\mathbf{x}^{(i)}; \boldsymbol{\theta})$ .

The posterior mean (21) serves as the cheap-to-evaluate surrogate for  $f(\mathbf{x})$  that we seek. Finally, the hyperparameters  $\boldsymbol{\theta}$  must be carefully chosen to ensure that the surrogate accurately represents  $f(\mathbf{x})$ . This is done by maximizing the marginal likelihood for the training data:

$$\log p(\mathbf{y}|\mathbf{X}) = -\frac{1}{2} \mathbf{y}^T [\mathbf{K} + \sigma_n^2 \mathbf{I}_N]^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K} + \sigma_n^2 \mathbf{I}_N| - \frac{N}{2} \log (2\pi). \quad (23)$$

Figure 2 illustrates a GPR in a one-dimensional setting in which the true function is  $f(x) = x \sin(x)$ . The example is for  $N = 5$  observations with observation noise  $\sigma_n = 0.01$ . The true function is modeled as a GP with specification (19) and (20) for the mean and covariance kernel, respectively. The posterior mean is calculated according to equation (21) and is shown in the solid line. The shaded region plots the 95% confidence interval around the posterior mean and is computed based on the posterior variance (22).

### 3.2 Deep Kernel Gaussian Process

We use the Deep Kernel Gaussian Process (DKGP) proposed by [Wilson, Hu, Salakhutdinov, and Xing \(2016\)](#). A DKGP is a GP whose covariance kernel embeds a neural network

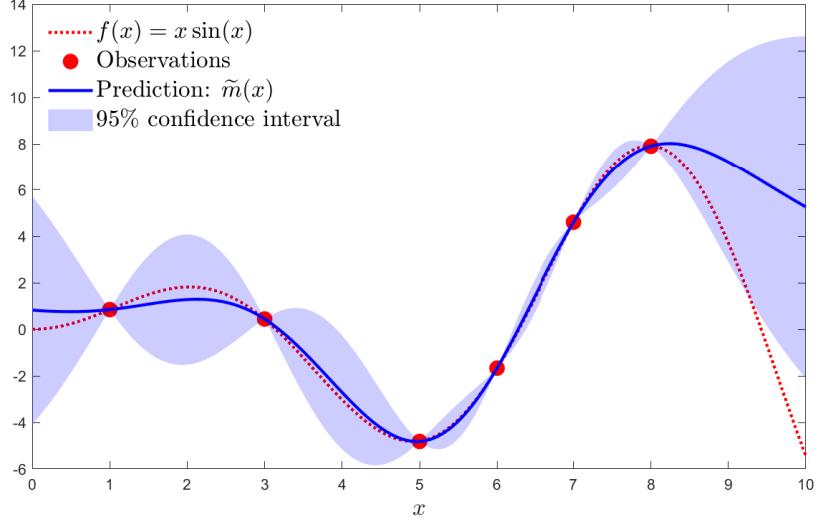


Figure 2: **Illustration of Gaussian Process Regression.** The dotted is the true function while the solid line is the posterior mean of the fitted GP. The shaded region plots the 95% confidence interval.

component within a “standard” kernel as follows:

$$k_{\text{DK}}(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}, \mathbf{w}) = k(NN(\mathbf{x}; \mathbf{w}), NN(\mathbf{x}'; \mathbf{w}); \boldsymbol{\theta}) \quad (24)$$

where  $NN(\cdot; \mathbf{w})$  is a neural network with hyperparameters  $\mathbf{w}$ , and  $k(\cdot, \cdot; \boldsymbol{\theta})$  is a “regular” kernel such as the Matern 5/2 kernel (20). That is, a deep kernel (24) first uses a neural network  $NN(\cdot; \mathbf{w})$  to process the input  $\mathbf{x}$  before feeding the extracted feature(s)  $NN(\mathbf{x}; \mathbf{w})$  into a standard kernel  $k(\cdot, \cdot; \boldsymbol{\theta})$ .

To see why it is necessary to use a deep kernel in our setting, consider Figure 3. Panel A illustrates a slice of the scaled value function  $v(t, w, k, \mu_P, s)$  along the liquid wealth share  $w$  dimension (holding all other variables fixed). We see that the value function consists of two relatively flat regions, below  $w = 0.15$  and above  $w = 0.30$ , with a sharp transition in the intermediate region between  $w = 0.15$  and  $w = 0.30$ .

The sharp transition in the value function is a generic feature of our problem and is due to the possibility of default. The logic is as follows. Defaulting next period is likely when current liquidity  $w$  is low; in this case, the current period’s value function primarily depends on the value of default  $v_D(t + 1, s')$  next period which is low. Conversely, defaulting next period becomes unlikely when current liquidity  $w$  is high and the current period’s value function primarily depends on the much larger non-default value next period  $v(t + 1, w', k', \mu'_P, s')$ . The sharp transition for intermediate liquidity  $w$  captures the transition between defaulting

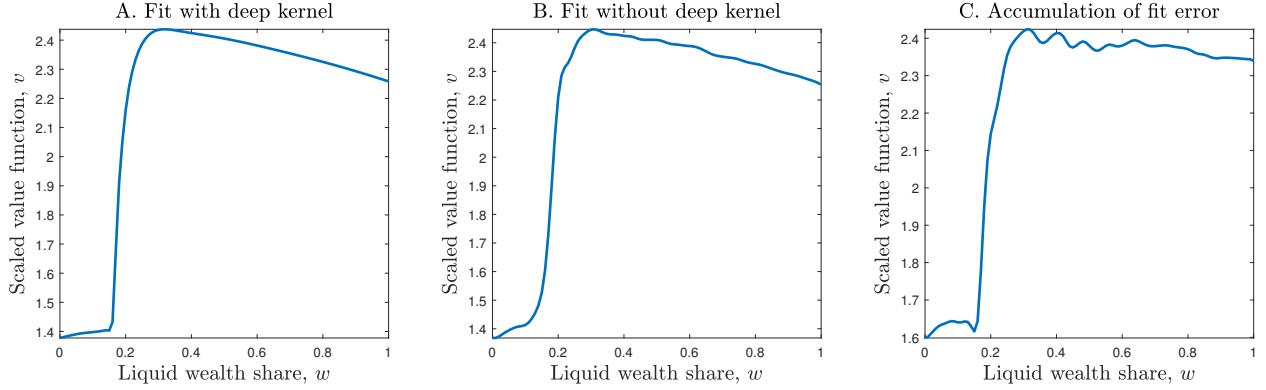


Figure 3: **Illustration of necessity of Deep Kernel Gaussian Process.** Panel A illustrates the value function when the model is solved using DKGPs at every single step. Panel B illustrates the solution from a single backward induction step when the value function is fitted using a standard kernel. Panel C illustrates the accumulation of error over successive steps of backward induction when the value function is fitted using a standard kernel.

and not defaulting next period.

Sharp transitions are challenging for standard kernels to capture. The reason is that standard kernels such as the Matern kernel (20) feature a constant lengthscale along each dimension. In contrast, the value function in our setting features several lengthscales—long lengthscales are necessary to capture the slowly-varying value function when  $w$  is low or high, while a short lengthscale is needed to capture the fast-varying value function for intermediate values of  $w$ . In addition, these lengthscales vary over the entire state space as both the location and speed of the sharp transition depend on all state variables.

Panels B and C of Figure 3 illustrates the pitfalls of using a single lengthscale to fit a multi-lengthscale value function. Maximum likelihood estimation fits a short lengthscale to capture the sharp transition. The resulting fitted GP is fast-varying and contains wiggles (see panel B). These wiggles accumulate over successive steps of backward induction and ultimately results in a very noisy value function after sufficiently many steps (see panel C of Figure 3).

The deep kernel (24) captures state-dependant lengthscales through its neural network component. This can be seen in Panel A of Figure 3 which plots the fitted value function when a DKGPs is used.

## 4 Quantitative Analysis

We begin by calibrating the model in Section 4.1 and describing the model’s solution in Section 4.2. We then report our model’s implications for optimal private asset allocation. Section 4.4 and Section 4.5 describe the optimal allocations over the LP’s life cycle and how business cycle conditions alter those allocations, respectively. Section 4.6 shows that, for long-term PE investors, whether or not PE returns must first be “unsmoothed” prior to making inferences and allocation decisions may be a moot point. Section 4.7 examines the impact of the risk weights on the LP’s portfolio allocation. Appendix C provides details on the numerical implementation of the model.

### 4.1 Calibration

We use the parameters listed in Table 1. We calibrated the model at a quarterly frequency and take the investment horizon to be  $T = 40$  quarters or 10 years.

We take state  $s = 1$  to be a recessionary state and state  $s = 2$  to be an expansionary state. We set the transition probabilities for the macroeconomic state to  $p_{12} = 0.25$  and  $p_{21} = 0.05$  based on the duration of NBER recessions in the post-WWII sample. This implies recessions and expansions last for 4 and 20 quarters on average, respectively. In what follows, we use NBER recession dates when estimating parameter values that depend on the macroeconomic state.

We use quarterly data provided by Liberty Mutual Investments (LMI) on its PE holdings to calibrate various PE-related parameters. This data contains LMI’s total PE exposure broken down by asset class—buyout, growth, and venture capital. For each asset class, the data details total exposure over time ( $K_t + P_t$  in the model) and breaks this exposure down into outstanding commitments ( $K_t$ ) and the NAV of current PE investments ( $P_t$ ). The data also include information on LMI’s contributions, distributions received from its PE investments, and the number of investees LMI invests in. In estimating parameters, we restrict attention to buyout funds from the moment when the number of buyout investees first reach 10. The resulting sample starts in 1996Q4 and ends in 2023Q4; all estimates are based on this sample period.

$\text{Contribution}_t$  in the LMI data corresponds to  $\lambda_K(s_t)K_{t-1} + \lambda_N(s_t)N_{t-1}$  in the model. We estimate the call rates using the following regression:

$$\frac{\text{Contribution}_t}{K_{t-1}} = a(s_t) + b(s_t)\frac{N_{t-1}}{K_{t-1}} + \epsilon_t, \quad (25)$$

Description	Symbol	Value	
		$s = 1$	$s = 2$
Investment horizon, quarters	$T$	40	
Macro transition probability	$p_{ss'}$	0.25	0.05
PE call rate, new commitments	$\lambda_N(s)$	0.18	0.047
PE call rate, existing commitments	$\lambda_K(s)$	0.050	0.078
PE distribution rate	$\lambda_D(s)$	0.028	0.071
PE liquidation discount	$\alpha(s)$	0.66	0.90
Risk-free rate	$\log R_f(s)$	0.0028	0.0051
Stock expected return	$\mu_S(s)$	0.0079	0.0238
Stock return volatility	$\sigma_S(s)$	0.1493	0.0829
Stock-PE return correlation	$\rho(s)$	0.9527	0.4575
PE return volatility	$\sigma_P(s)$	0.0768	0.0424
PE expected return autocorrelation	$\varrho_{P,1} + \varrho_{P,2}$		0.201
PE expected return state-dependance	$\nu_P(s)$	0.0024	0.0317
Risk aversion	$\gamma$	2	
PE adjustment costs	$\gamma_N$	0.1	
PE adjustment costs	$\bar{n}$	0	
Stock adjustment costs	$\gamma_S$	0.01	
Risk budget threshold	$\bar{\theta}$	1	
Risk budget cost	$\kappa$	1	
Risk weight, bonds	$\theta_B$	0	
Risk weight, stocks	$\theta_S$	1.5	
Risk weight, PE	$\theta_P$	1.5	

Table 1: **Parameters.** We simulate the model at a quarterly frequency using the parameters in this table.

where the constant  $a(s_t)$  maps to  $\lambda_K(s_t)$  and the slope coefficient  $b(s_t)$  maps  $\lambda_N(s_t)$ . We estimate regression (25) separately for recessionary ( $s = 1$ ) and expansionary ( $s = 2$ ) regimes to obtain state-dependant estimates for the call rates. The resulting point estimates give  $\lambda_K(1) = 0.050$  and  $\lambda_N(1) = 0.18$  during recessions, and  $\lambda_K(2) = 0.078$  and  $\lambda_N(2) = 0.047$  during expansions.

We set the PE payout rate  $\lambda_D(s)$  to be the average payout rate in the LMI dataset. This results in  $\lambda_D(1) = 0.028$  and  $\lambda_D(2) = 0.071$  during recessions and expansions, respectively.

We calibrate the discount from liquidating PE investments  $\alpha(s)$  based on estimates of transaction costs in the secondary market for selling PE stakes from [Nadauld, Sensoy, Vorkink, and Weisbach \(2019, Table 1\)](#). They find that the purchase price, expressed as a percent of NAV, is 86% on average and drops to 66% during 2008-09 (the only recession in their sample). We set  $\alpha(1) = 0.66$  and  $\alpha(2) = 0.90$  during recessions and expansions, respectively, based on

these estimates.

We calibrate the returns processes as follows. We take the series for the risk-free rate and aggregate stock returns from Kenneth French's website. We set the risk-free rate and expected stock returns to their sample averages during recessions and expansions; this results in  $\log R_f(1) = 0.0028$  and  $\log R_f(2) = 0.0051$  for the risk-free rate and  $\mu_S(1) = 0.0079$  and  $\mu_S(2) = 0.0238$  for expected stock returns.

LMI data includes quarterly mark-to-market information for the NAV of its PE investments; this allows us to compute the realized PE return  $R_{P,t}$  as the ratio of the mark-to-market and the lagged NAV value. We use moments of the resulting time series for  $R_{P,t}$  to estimate the PE return process. We set the stock-PE return correlation  $\rho(s)$  to its sample correlation; this results in  $\rho(1) = 0.9527$  and  $\rho(2) = 0.4575$  during recessions and expansions, respectively. We restrict  $\varrho_{P,1} = \varrho_{P,2}$  and choose the remaining parameters of the PE expected return process (13) and PE return volatility  $\sigma_P(s)$  by matching the following data moments: (1) a PE expected return  $\mathbb{E}[\log R_{P,t+1}|s_t]$  of 0.0052 and 0.0392 during recessions and expansions, respectively, (2) a PE return volatility  $\sigma(\log R_{P,t+1}|s_t)$  of 0.0772 and 0.0427 during recessions and expansions, respectively, and (3) a PE return autocorrelation of  $\rho(R_{P,t}, R_{P,t+1}) = 0.1425$ . This results in  $\varrho_{P,1} = \varrho_{P,2} = 0.1006$ ,  $\nu_P(1) = 0.0024$  and  $\sigma_P(1) = 0.0768$  during recessions, and  $\nu_P(2) = 0.0317$  and  $\sigma_P(2) = 0.0424$  during expansions.

We set the risk weights for stocks and private equity to  $\theta_S = 1.5$  and  $\theta_P = 1.5$ , respectively. Since the overall risk budget  $\bar{\theta}$  is normalized to 1, these risk weights imply a 50% risk charge for stocks and private equity. These values are in line with equity risk charges used by ratings agencies for assessing insurers' RBC adequacy (see [S&P Global 2023](#), Table 14). The risk weight for risk-free bonds is  $\theta_B = 0$ . We set the cost parameter for the proportional risk cost (6) to  $\kappa = 1$ .

We set risk aversion to  $\gamma = 2$  which is standard. The remaining parameters relate to adjustment costs. We set  $\gamma_N = 0.1$  and  $\bar{n}$  for PE adjustment costs, and  $\gamma_S = 0.01$  for stock adjustment costs.

## 4.2 Value and policy functions

We begin by illustrating the solution for the scaled value function and the policy functions; recall that the scaled and unscaled value functions are related through equations (17) and (18).

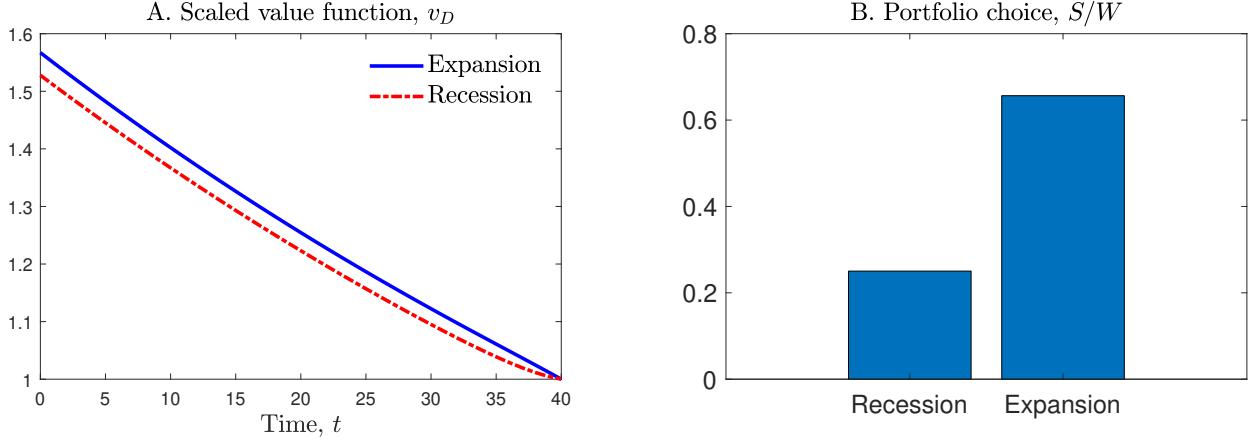


Figure 4: **Solution after defaulting.** Panel A plots the scaled value function after defaulting  $v_D(t, s)$ ; it is related to the unscaled value function through equation (17). Panel B plots the optimal allocation to stocks after defaulting.

**Value functions.** Figure 4 plots the solution after defaulting. Panel A plots the scaled value function  $v_D(t, s)$ . The value function is declining over time since the investment horizon over which to accumulate wealth shrinks over time. The value function is also higher during expansions than recessions. Panel B plots the portfolio choice after defaulting. The investor’s allocation depends only on the macroeconomic state and does not vary over time—the optimal allocation to stocks is 25% and 66% during recessions and expansions, respectively; all remaining wealth are allocated to bonds.

Next, we illustrate the solution before defaulting during which the scaled value function  $v(t, w, k, \mu_P, s)$  depends on time  $t$ , the liquid fraction of total wealth  $w = W/(W + P)$ , uncalled commitments relative to total wealth  $k = K/(W + P)$ , the expected return for PE  $\mu_P$ , and the macroeconomic state  $s$ .

Figure 5 plots the distribution of the expected PE return  $\mu_P$ . Panel A plots the unconditional distribution while panel B plots the conditional distribution given the macroeconomic state. Expected returns for PE are higher during expansions under our baseline calibration with expected returns averaging  $\mathbb{E}[\mu_{P,t}|s_t = 1] = 0.0052$  and  $\mathbb{E}[\mu_{P,t}|s_t = 2] = 0.0392$  per quarter during recessions and expansions, respectively. From panel B, we also see that the distribution of expected PE returns is wider during expansions than recessions; this is on account of higher PE return volatilities during recessions.

Figure 6a illustrates the scaled value function  $v(t, w, k, \mu_P, s)$  at  $t = 0$ . Panel A plots the value function in the expansionary state ( $s = 2$ ) when the expected return on PE is equal to its mean conditional on  $s = 2$ . We see that the value function  $v$  is a non-monotonic function of the liquid fraction of total wealth  $w$ . Specifically,  $v$  is first increasing in  $w$  when

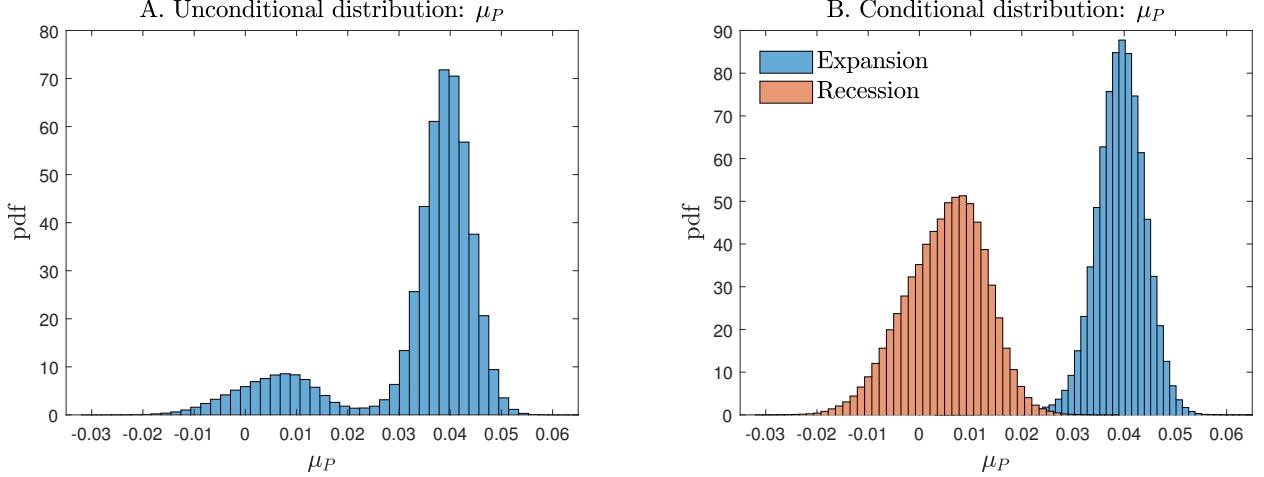


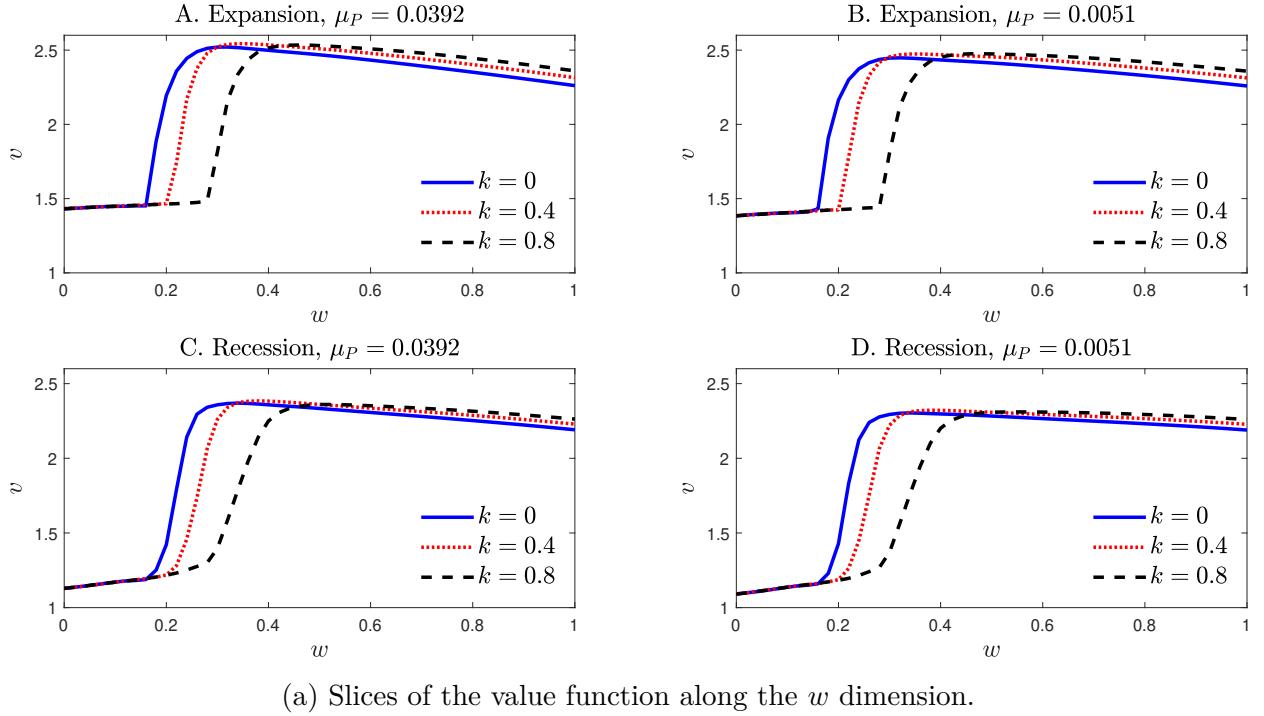
Figure 5: **Distribution for  $\mu_P$ .** Panel A plots the stationary distribution for  $\mu_P$  whose law of motion is given by equation (13). Panel B plots the distribution for  $\mu_P$  conditional on the current macroeconomic state.

$w$  is sufficiently small (e.g., for  $w < 0.30$  when  $k = 0$ ) and then decreasing in  $w$  when  $w$  is sufficiently large (e.g., for  $w > 0.30$  when  $k = 0$ ).

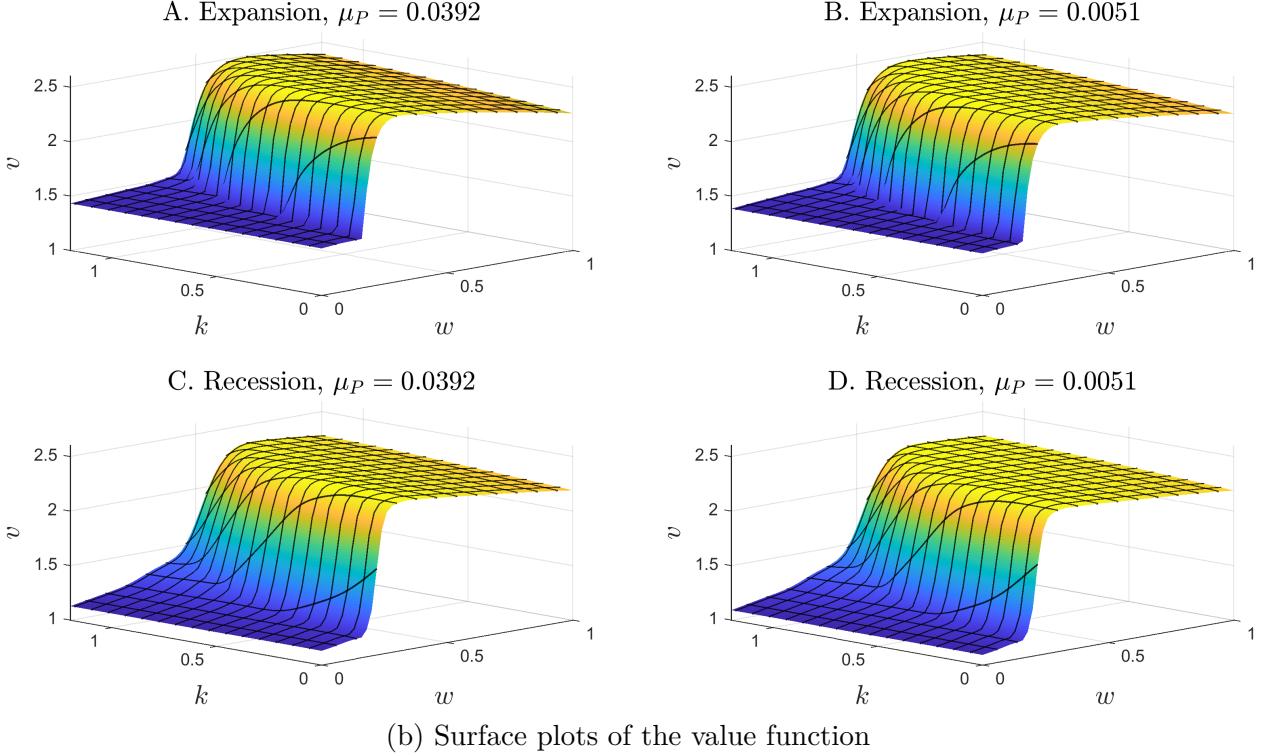
The reason for this non-monotonicity is as follows. Defaulting becomes likely when  $w$  is small; this is illustrated in Figure 7 which plots the one quarter ahead default probabilities at  $t = 0$ . When the current period's liquid wealth is low, the LP becomes less likely to have enough liquid wealth next period to meet its capital commitments and to pay its risk budget cost (see equation (9)). Hence, the value function becomes increasing in  $w$  when  $w$  is small, as having more liquidity helps the LP avoid default. Instead, when  $w$  is sufficiently large so that there is no risk of default, the LP's value function becomes decreasing in  $w$  because the LP would be able to achieve better investment outcomes had it allocated more of its wealth to PE.

Similarly, the value function  $v$  is non-monotonic in uncalled commitments relative to total wealth  $k$ . For example, in panel A of Figure 6a we see that  $v$  is decreasing in  $k$  for small  $w$  (e.g., when  $w = 0.2$ ) while  $v$  is increasing in  $k$  for large  $w$  (e.g., when  $w = 0.8$ ). The reason is as follows. When the LP has little liquidity, default concerns become first order and having more uncalled commitments increases the likelihood of default, thereby decreasing  $v$ . In contrast, when the LP has sufficient liquidity and default is no longer a concern, higher uncalled commitments  $k$  increases the LP's PE exposure which increases  $v$ . This is because the presence of adjustment costs means that it takes time for the LP to build up its PE exposure through capital commitments.

Panel B of Figure 6a plots the value function in the expansionary state when the PE



(a) Slices of the value function along the  $w$  dimension.



(b) Surface plots of the value function

Figure 6: **Scaled value function at  $t = 0$ .** This figure illustrates the scaled value function before default  $v(t, w, k, \mu_P, s)$  at  $t = 0$ . In both subfigures, the first (second) row displays the value function in the expansionary (recessionary) state. The first (second) column sets  $\mu_P$  to its mean conditional on the expansionary (recessionary) state.

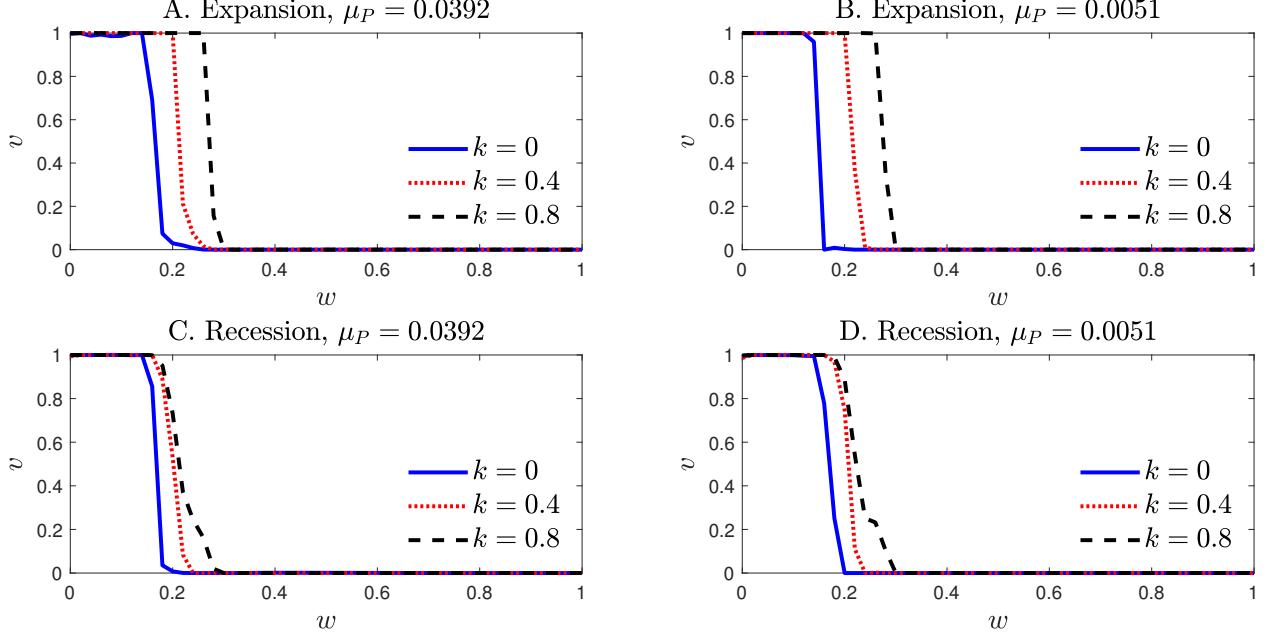


Figure 7: **One quarter ahead default probabilities at  $t = 0$ .** This figure illustrates the one period ahead default probability at  $t = 0$ . The first (second) row displays the default probabilities in the expansionary (recessionary) state. The first (second) column sets  $\mu_P$  to its mean conditional on the expansionary (recessionary) state.

expected return is low while panels C and D illustrate the value function in the recessionary state. The value function is monotonically increasing in the PE expected return  $\mu_P$ , as higher expected PE returns unambiguously translate into higher expected utility. Moreover, the value function is lower during recessions (panels C and D) compared to expansions (panels A and B), reflecting the less favorable investment opportunities during economic downturns. Despite these level differences, we see that the shape of the value function remains similar across different macroeconomic states and PE expected returns. Figure 6b illustrates further details of the value function using three-dimensional surface plots.

**New commitments.** Figure 8 illustrate the optimal policy for new PE commitments, scaled with respect to total wealth, at  $t = 0$ . We see that new PE commitments are increasing in the liquid fraction of total wealth  $w$  and decreasing in the uncalled commitments relative to total wealth  $k$ . The LP becomes less willing to commit new capital to PE when default concerns are high as defaulting leads to costly PE liquidations. This is why new PE commitments are low when  $w$  is low or  $k$  is high. Instead, when default is unlikely, the LP commits more capital to PE in order to increase its PE exposure and thereby lower its future liquid wealth share (as discussed above, the value function becomes decreasing in  $w$  in the absence of default concerns).

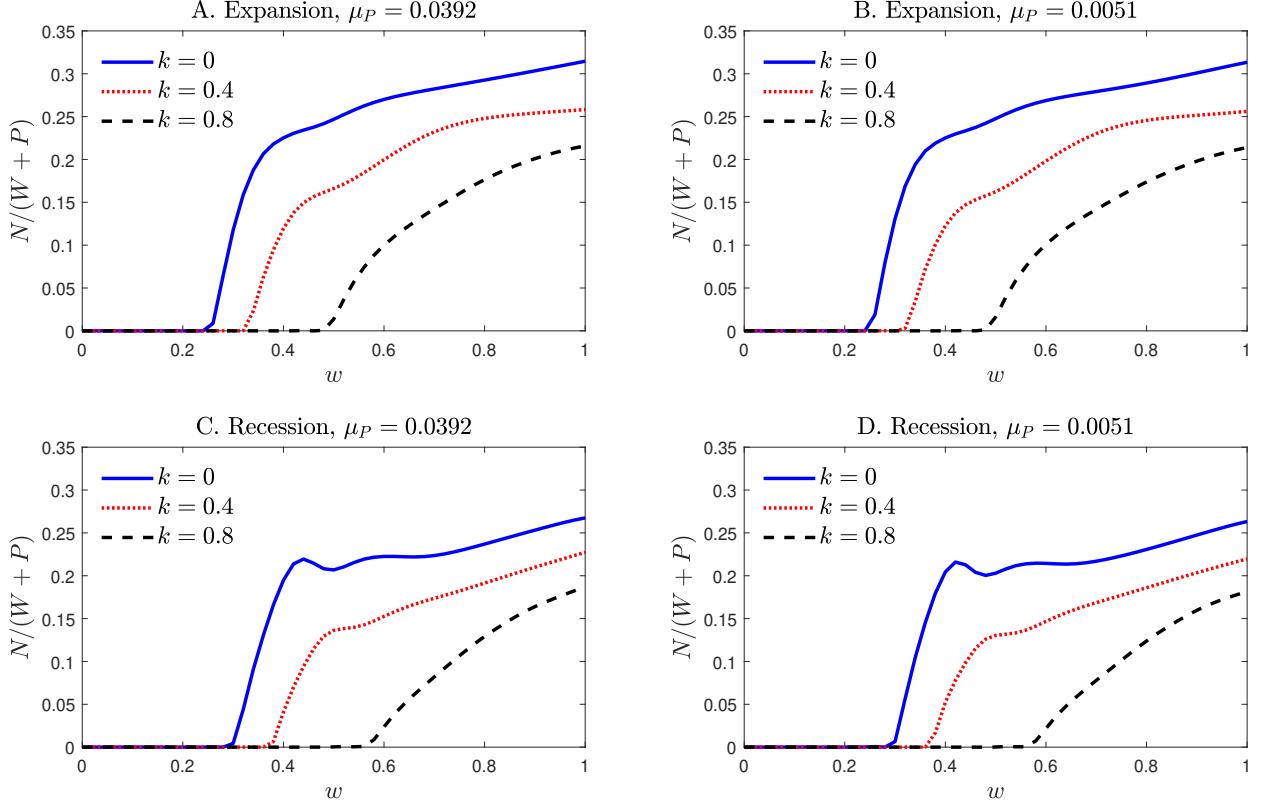


Figure 8: **New commitments at  $t = 0$ .** This figure illustrates new commitments relative to total wealth,  $N/(W + P)$ , at  $t = 0$ . The first (second) row displays the results for the expansionary (recessionary) state. The first (second) column sets  $\mu_P$  to its mean conditional on the expansionary (recessionary) state.

Comparing panels A and B of Figure 8, and panels C and D of Figure 8, we see that current PE expected returns  $\mu_P$  have surprisingly little impact on new commitments across both expansionary and recessionary states. This muted response stems from the inherent delay between making commitments and their eventual conversion into PE investments. Because PE expected returns exhibit low quarterly autocorrelation in the baseline calibration ( $\varrho_P = 0.20$ ), initial differences in  $\mu_P$  largely dissipate before committed capital is called. Consequently, optimal commitment levels depend primarily on the expected conditional return  $\mathbb{E}[\mu_{P,t}|s_t]$  rather than the current level of  $\mu_P$ , since only invested capital (NAV) ultimately generates PE returns.

Finally, comparing across macroeconomic states, we see from Figure 8 that new commitments are higher during expansions than recessions. Figure C.1a in Appendix C uses three-dimensional surface plots to illustrate further details regarding new commitments.

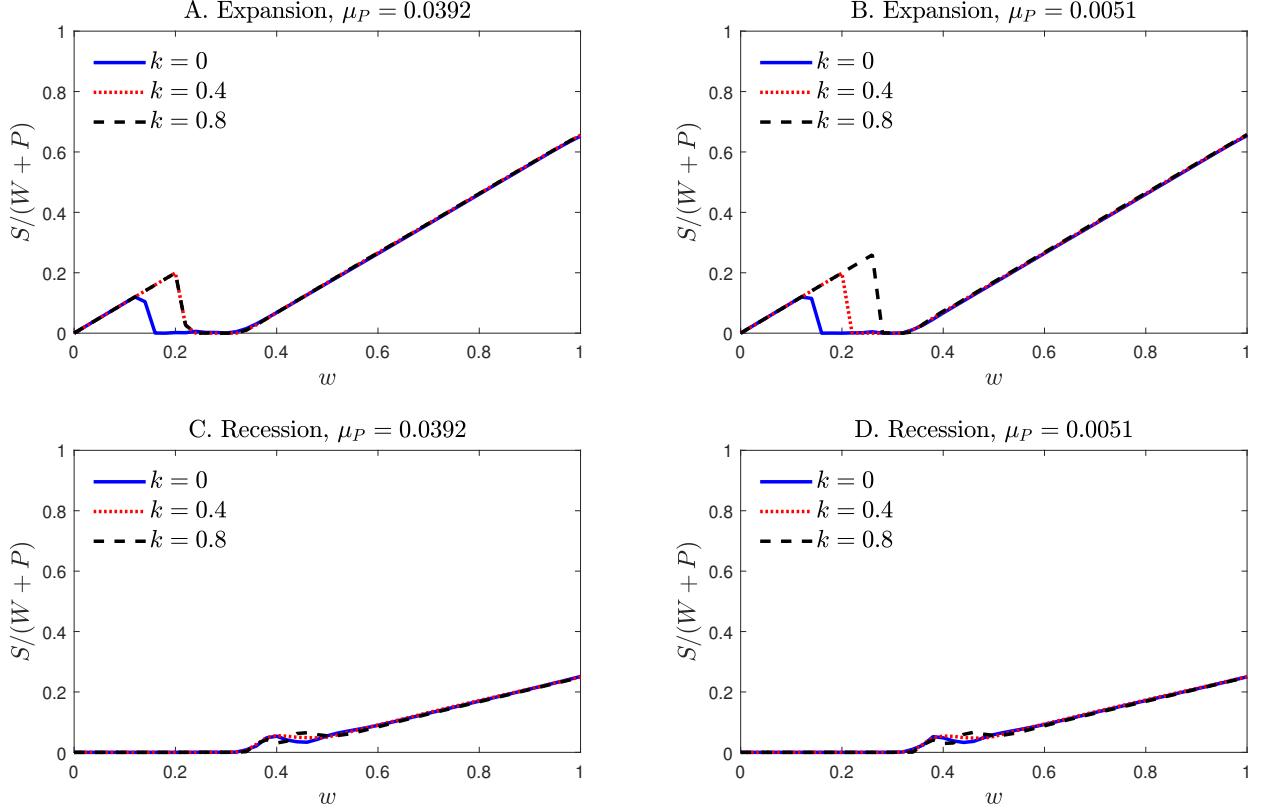


Figure 9: **Allocation to stocks at  $t = 0$ .** This figure illustrates the optimal allocation to public stocks, expressed as a fraction of total wealth, at  $t = 0$ . The first (second) row displays the optimal allocation in the expansionary (recessionary) state. The first (second) column sets  $\mu_P$  to its mean conditional on the expansionary (recessionary) state.

**Allocation to stocks.** Figure 9 illustrates the optimal allocation to public stocks as a fraction of total wealth at  $t = 0$  (see Figure C.1b in Appendix C for surface plots that display further details). Panel A depicts this allocation during the expansionary state, with PE expected returns at their mean conditional on  $s = 2$ . When liquid wealth  $w$  exceeds 0.3, making default risk negligible, two patterns emerge. First, stock allocation increases with  $w$ , reflecting a rebalancing mechanism: as liquid wealth rises and PE exposure  $(1 - w)$  mechanically falls, the LP maintains its desired risk exposure by increasing its allocation to public stocks. For example, the portfolio shifts from 50% PE, 17% stocks, and 33% bonds at  $w = 0.5$  to 20% PE, 46% stocks, and 34% bonds at  $w = 0.8$ . Second, stock allocation remains independent of uncalled commitments  $k$ . This is because, absent default risk, the value function depends primarily on next period's total wealth growth rate. Since uncalled commitments affect total wealth only after their conversion to PE investments (NAV), which occurs with a lag, they have no immediate impact on portfolio growth (see equation (B.9)) and therefore do not interact with the allocation decision for public stocks.

When default risk is non-negligible, a different mechanism is at play. For example, default occurs with certainty when  $w = 0.1$  for the levels of  $k$  considered here (see Figure 7). In this case, the LP makes its stock allocation decision knowing that it will default on its PE commitments and incur a PE liquidation loss next period. The LP chooses to maximize its stock allocation to compensate for this loss when macroeconomic conditions are favorable (see panels A and B of Figure 9). Conversely, it takes a conservative approach and does not allocate anything to public stocks when macroeconomic conditions are unfavorable (see panels C and D of Figure 9).

However, this default mechanism is rarely at work along the optimal path when we simulate the model. The LP tries to avoid default since defaulting is costly. The cumulative default rate over the LP's 10 year investment horizon is only 0.1%.

### 4.3 Marginal value of liquidity

The illiquid nature of PE assets makes liquidity management central. Indeed, the LP must finance its capital calls out of its liquid wealth or suffer the consequences of default. We use our model to quantify the value of liquidity through two measures.

The first measure, the marginal value of liquidity at time  $t$ , measured in units of certainty equivalent terminal wealth, is given by

$$\frac{\partial}{\partial W}(W + P)v(t, w, k, \mu_P, s) = v + (1 - w)v_w - kv_k. \quad (26)$$

On the left-hand side of equation (26),  $(W + P)v(t, w, k, \mu_P, s)$  expresses the value function  $V(t, w, k, \mu_P, s)$  in certainty equivalent units of terminal wealth (see equation (18)). Here,  $W + P$  is current total wealth and  $v(t, w, k, \mu_P, s)$  is the certainty equivalent value per unit of wealth. On the right-hand side, the  $v$  term captures the direct effect of an additional unit of wealth; the  $(1 - w)v_w$  and  $-kv_k$  account for compositional effects due to changes in the liquid fraction of wealth  $w$  and uncalled commitments relative to total wealth  $k$ . The formulation of the marginal value of liquidity (26) is analogous to the marginal value of cash in the liquidity management model of [Bolton, Chen, and Wang \(2011, 2013\)](#).

What forces determine the marginal value of liquidity (26)? In bad times, liquidity could be used to stave off default. In good times, the LP prefers to convert liquidity to high yielding illiquid private assets. However, that will take time since due to the lag between making commitments and getting capital called, after which the return becomes the same as that of an already existing unit of illiquid wealth that has already been put to use.

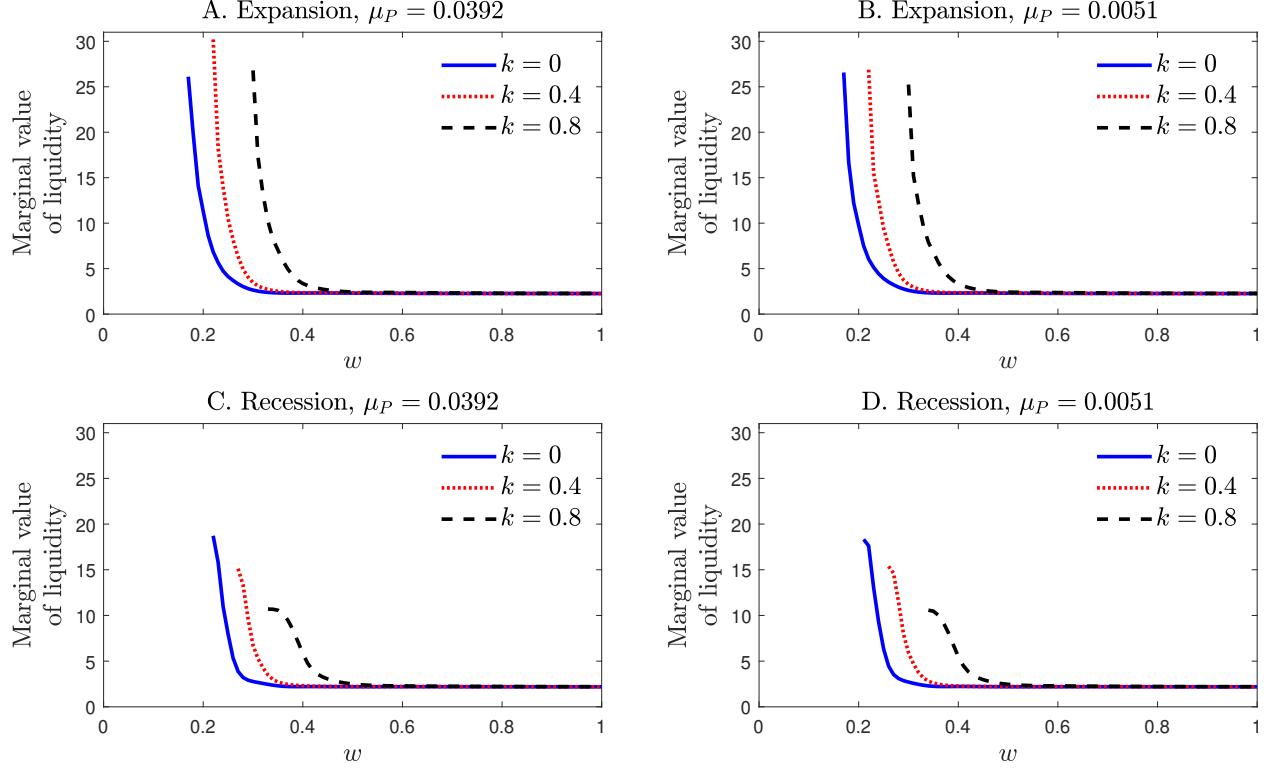


Figure 10: **Marginal value of liquid wealth at  $t = 0$ .** This figure illustrates the marginal value of liquid wealth (26) at  $t = 0$ . The first and second rows display the marginal value of cash in the expansionary and recessionary states, respectively; the first and second columns set  $\mu_P$  to its mean conditional on the expansionary and recessionary state, respectively.

Figure 10 illustrates the marginal value of liquid wealth (26) at  $t = 0$ . We see that the marginal value of liquid wealth is highly sensitive to default risk. As the liquid fraction of wealth  $w$  decreases, default probabilities increase (see Figure 7), leading to a spike in the marginal value of liquidity. For instance, panel A of Figure 10 shows that when  $k = 0.4$ , the macroeconomic state is expansionary, and  $\mu_P$  equals its conditional mean in the expansionary state, the marginal value of liquidity equals 2.28 when  $w = 1$  and default is not a concern. However, as  $w$  drops to 0.22, the marginal value of liquidity rises to 30.21; at this point, the one-quarter-ahead default probability is 0.21 and rapidly approaches 1 as  $w$  declines further. Figure 10 further demonstrate how the marginal value of liquidity varies over the state space. For example, the location of the peak in the marginal value of liquidity depends strongly on the ratio of uncalled capital to wealth  $k$ —higher values of uncalled capital  $k$  increases the marginal value of liquidity (26), especially at low levels of liquidity  $w$ .

The second measure of the value of liquidity is the [Ang, Papanikolaou, and Westerfield](#)

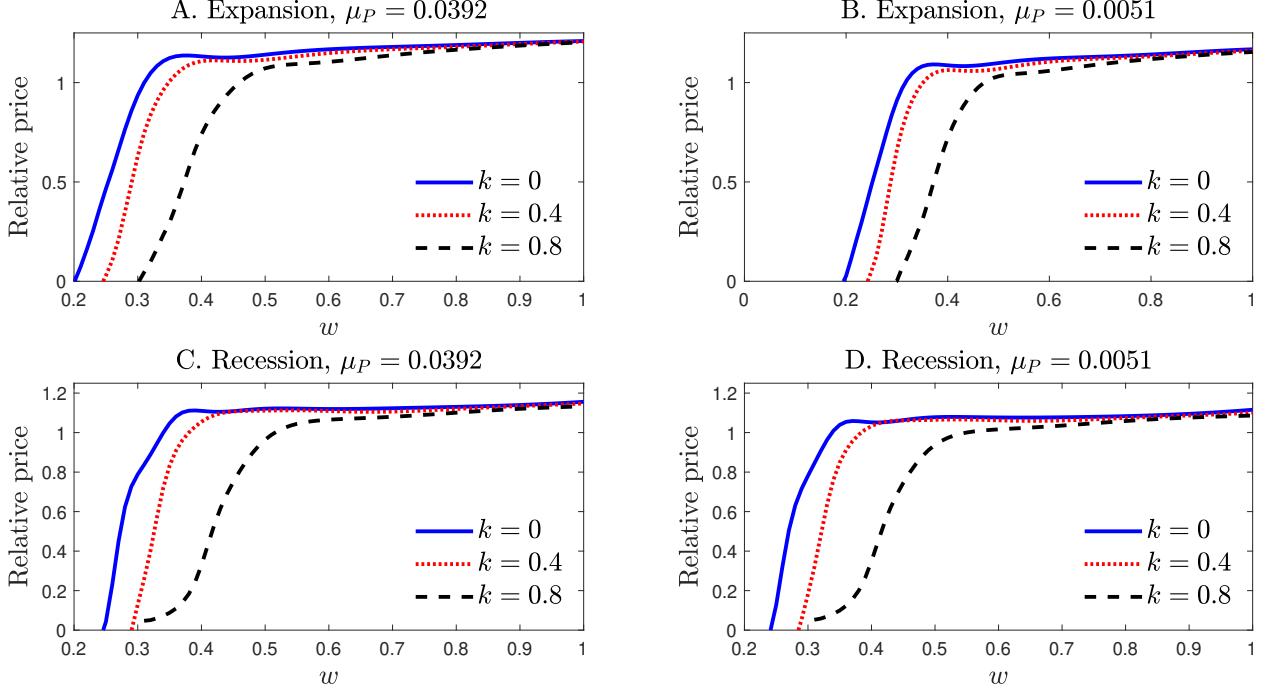


Figure 11: **Relative price of illiquid to liquid wealth at  $t = 0$ .** This figure illustrates the relative price of illiquid to liquid wealth (27) at  $t = 0$ . The first and second rows display this relative price in the expansionary and recessionary states, respectively; the first and second columns set  $\mu_P$  to its mean conditional on the expansionary and recessionary state, respectively.

(2014) measure of liquidity, which becomes

$$q = \frac{\frac{\partial}{\partial P}(W + P)v(t, w, k, \mu_P, s)}{\frac{\partial}{\partial W}(W + P)v(t, w, k, \mu_P, s)} \quad (27)$$

when adapted to our setting. The denominator in equation (27) is already computed in equation (26) while the numerator equals  $\frac{\partial}{\partial P}(W + P)v(t, w, k, \mu_P, s) = v - wv_w - kv_k$ . The ratio (27) is interpreted as a relative price ratio: the value of an additional unit of illiquid wealth measured in terms of liquid wealth.

Figure 11 illustrates the relative price ratio  $q$  at  $t = 0$ . We see that  $q$  is highly sensitive to default risk;  $q > 1$  when there is sufficient liquid wealth  $w$  to stave off default risk so that high yielding illiquid wealth can mature without the worry of default. When  $w$  is low,  $q < 1$  due to the possibility of liquidating illiquid wealth at a loss in case of default. These features are in line with the findings in [Ang, Papanikolaou, and Westerfield \(2014\)](#). We build on [Ang, Papanikolaou, and Westerfield \(2014\)](#) by quantifying the nonlinear dependance of  $q$  on uncalled capital  $k$ , macroeconomic conditions  $s$ , and variation in expected PE returns  $\mu_P$ . Specifically, we see that  $q$  is decreasing in  $k$ , increasing in  $\mu_P$ , and is higher during expansions.

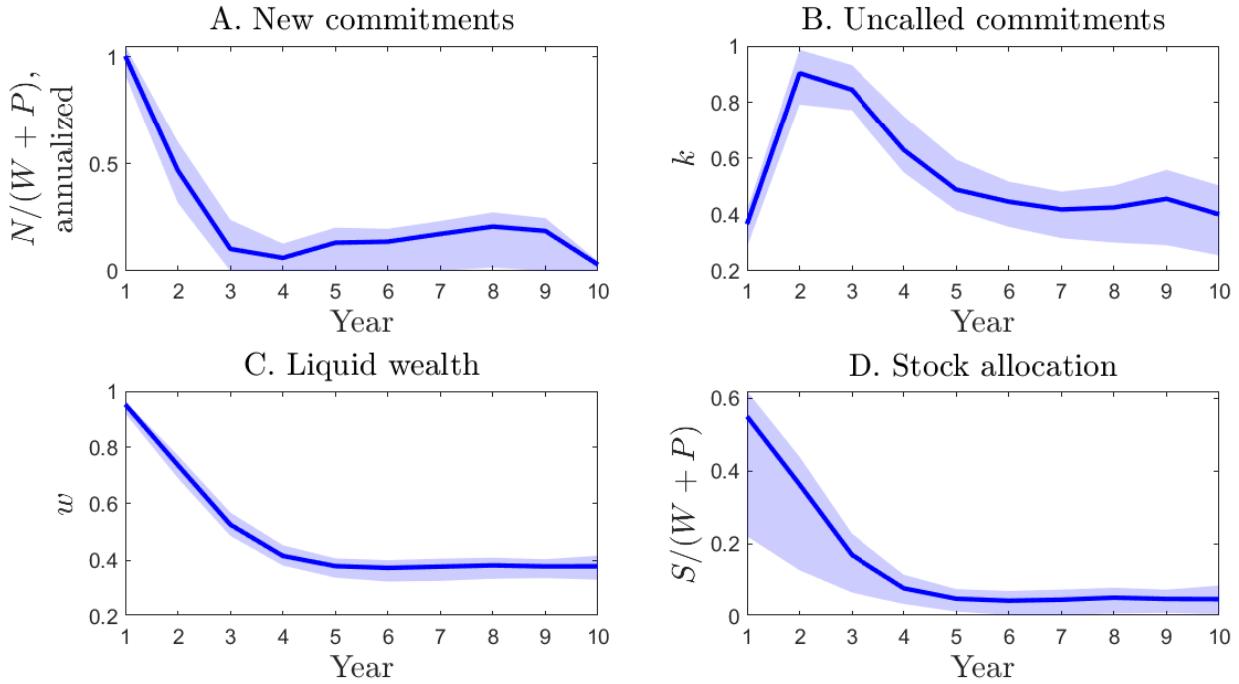


Figure 12: **Life cycle dynamics.** This figure illustrates the life cycle dynamics of the LP’s portfolio. The solid line plots the average outcome in each year while the shaded region provides the 95% confidence interval. Panels A through D plot the outcomes for new commitments, uncalled commitments, liquid wealth, and the stock allocation, respectively. All outcomes are shown relative to total wealth.

There are also nonlinear interaction effects. For example,  $q$  decreases quickly as  $k$  increases for low  $w$ , while this effect is attenuated as  $w$  increases.

#### 4.4 Life cycle dynamics

We now turn to the life cycle dynamics of the LP’s portfolio. We assume that the LP starts out at  $t = 0$  without any PE investments (i.e.,  $P_0 = K_0 = 0$ ) and  $w_0$  units of liquid wealth  $W_0 = 1$  so that  $k_0 = 0$  and  $w_0 = 1$ . The initial macroeconomic state  $s_0$  and PE expected return  $\mu_{P,0}$  are drawn from their stationary distributions. We then simulate the LP’s portfolio allocation over its investment horizon of 10 years. Figure 12 shows the outcome—the solid line plots the average outcome at each point in time while the shaded region plots the 95% confidence interval around the average.

Starting out, the LP aggressively makes new capital commitments in the first two years to rapidly increase its PE exposure (see panel A of Figure 12). For example, new commitments average 100% of total wealth during the first year before dropping to 47% and 10% of total

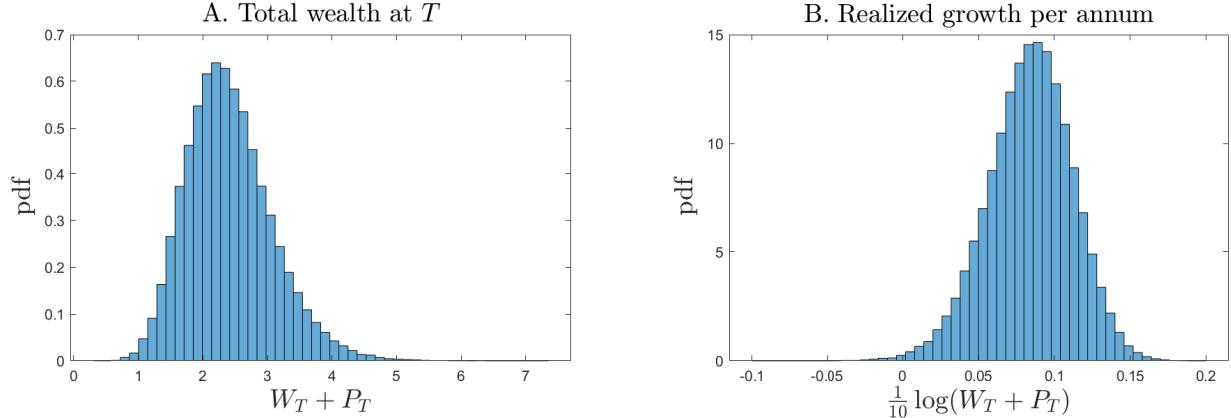


Figure 13: **Total wealth outcome distribution.** This figure illustrates the distribution of total wealth outcomes at the end of the investment horizon  $t = T$ . The initial total wealth is  $W_0 + P_0 = 1$ . Panel A plots the distribution for total wealth  $W_T + P_T$ . Panel B plots the distribution for the realized growth in total wealth per annum,  $\log(W_T + P_T)/10$ .

wealth on average in the second and third years, respectively. As a result, uncalled capital increases and peaks at an average of 90% and 85% of total wealth during the second and third years (see panel B). This is accompanied by a drop in the liquid portion of wealth (see panel C) or, equivalently, an increase in the LP's PE exposure. Specifically, the LP's NAV in PE investments rises to 26% and 47% of total wealth on average in the second and third years, respectively, before reaching 62% of total wealth by year 5. From year 5 onwards, the LP is in maintenance mode, making new commitments to maintain its PE exposure at an average of 62% of total wealth until the investment horizon is reached at year 10.

Stock allocation initially starts high (see panel D); for example, stock allocation averages 55% of total wealth in the first year. This is because the LP can only achieve its target aggregate risk exposure through public stocks before its PE exposure is built up. The wide confidence interval for stock allocation during this initial phase is due to differences in stock allocation across the business cycle which we discuss in Section 4.5. The stock allocation subsequently drops to 5% of total wealth on average during the LP's maintenance phase.

Figure 13 shows the distribution of outcomes at the terminal date  $T$ . Panel A plots the distribution of for total wealth  $W_T + P_T$ . Starting from a unit initial total wealth, the LP's terminal wealth equals 2.41 on average with a standard deviation of 0.66. This wealth distribution translates into a certainty equivalent wealth of  $\mathbb{E}[v(t = 0, w_0, k_0, \mu_P, s) | w_0 = 1, k_0 = 0] = 2.25$  on the initial date. Panel B illustrates the distribution of final outcomes in terms of the annualized realized total returns  $\frac{1}{10} \log((W_T + P_T)/(W_0 + P_0))$ . The annualized realized total return over the LP's investment horizon has a mean of 8.41% and a standard deviation of 2.78%. Note that this seemingly low standard deviation is because we are reporting the

standard deviation of annualized long-horizon returns (see Footnote 2).

#### 4.4.1 Heuristic versus dynamic private asset allocation

In practice, investors often use heuristics to simplify private asset allocation decisions. One common heuristic is to split the allocation problem into two steps. In the first step, investors determine a long-run target portfolio by solving a static portfolio choice problem that abstracts from liquidity risk and the timing lags inherent in PE investing (e.g., a traditional Markowitz approach). In the second step, they adjust their allocations to reach the targets from the first step by accounting for the delays between capital commitments, capital calls, and eventual distributions (see, for example, [Takahashi and Alexander 2002](#)). We demonstrate that portfolio allocations derived from our dynamic model can differ dramatically from these heuristic approaches. Specifically, we show that portfolio outcomes from our dynamic model, starting from year 5 when the LP has reached the maintenance phase, substantially differ from the long-run target portfolio obtained under the heuristic approach. We provide details below.

Within our framework, the first step of the heuristic approach corresponds to solving

$$V_{\text{heuristic}}(t, W_{\text{total}}, s) = \max_{S, B, P \geq 0} \mathbb{E} [V_{\text{heuristic}}(t + 1, W'_{\text{total}}, s') | W_{\text{total}}, s] \quad (28)$$

subject to

$$\begin{aligned} W'_{\text{total}} &= R'_P P R'_S S + R_f(s) B - W_{\text{total}} \Gamma(\theta), \\ \theta &= \frac{\theta_B B + \theta_S S + \theta_P P}{B + S + P}, \\ W_{\text{total}} &= S + B + P, \end{aligned}$$

and the terminal condition  $V_{\text{heuristic}}(T, W, s) = W^{1-\gamma}/(1-\gamma)$ . The heuristic problem (28) resembles the full problem (16) but reduces the total number of state variables by abstracting away from the illiquidity of PE investments. In implementing the heuristic problem (28), we further set  $\mu_P = \mathbb{E}[\mu_{P,t} | s_t = s]$  to its mean conditional on the macroeconomic state. In [Appendix B](#), we show that the solution to the heuristic problem (28) amounts to optimizing the certainty-equivalent growth rate of total wealth at each point in time. The solution is static in the sense that the optimal portfolio weights depend solely on the macroeconomic state  $s$  and do not otherwise change over time.

Figure 14 compares the portfolio outcomes under the heuristic approach (28) with those under our full dynamic model (16). The circles represent the portfolio allocations derived

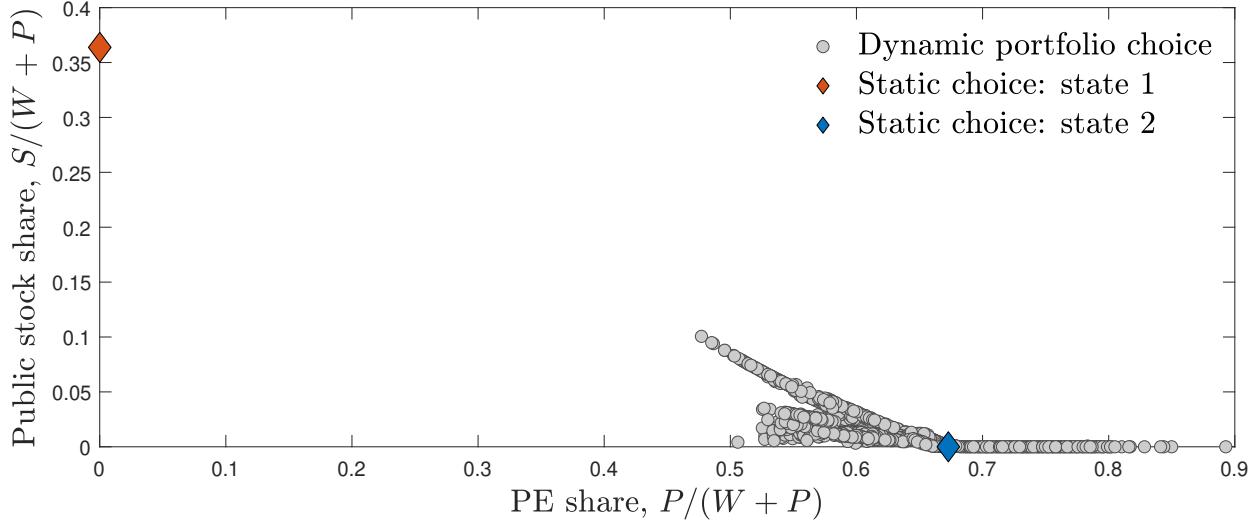


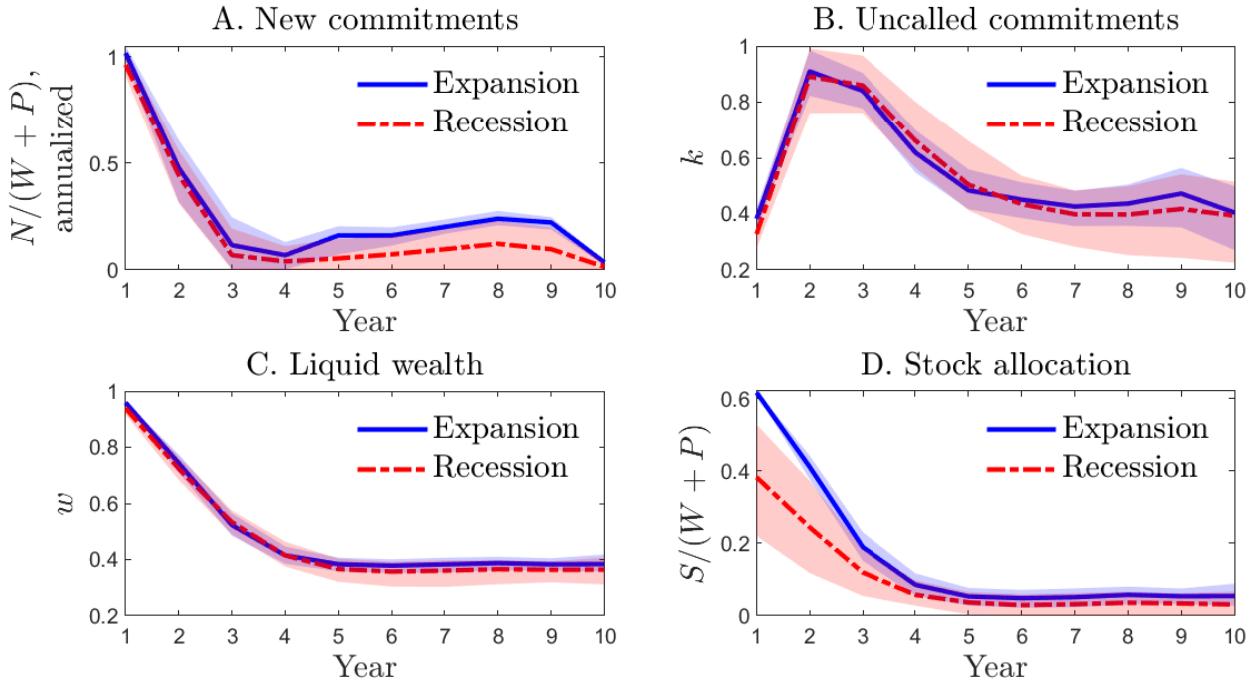
Figure 14: **Portfolio choice: heuristic vs. dynamic private asset allocation.** This figure compares the portfolio outcomes under the heuristic problem (28) (shown by the diamonds) against the outcomes under the full dynamic allocation problem (16) (shown in the circles). Outcomes for the latter correspond are from year 5 onwards during which the maintenance phase has been reached.

from our dynamic model for year 5 onwards—when the limited partner (LP) has entered its maintenance phase—while the diamonds show the static portfolio choices obtained from solving (28) for each macroeconomic state.

Under our baseline calibration (Table 1), the heuristic approach yields an optimal portfolio allocation of 67.3% in private equity (PE) and 0% in stocks (with the remainder in bonds) in the expansionary state, and 0% in PE and 36.4% in stocks in the recessionary state. In contrast, the optimal dynamic allocation maintains a PE allocation of at least 45% during the maintenance phase, while the stock allocation does not exceed 10%. This difference in outcomes is especially stark during the recessionary state, where the heuristic outcome falls entirely outside the range of values generated by the optimal dynamic allocation.

## 4.5 Business cycle dynamics

Figure 15 illustrates how business cycle conditions affect the life cycle dynamics of the LP’s portfolio. The solid and dashed lines plot the average outcomes conditional on the macroeconomic state being in a expansion and a recession, respectively; the shaded 95% confidence intervals are also conditional on the macroeconomic state. Overall, we see that the allocation patterns across the business cycle are similar compared to the unconditional case discussed in Section 4.4.



**Figure 15: Life cycle dynamics over the business cycle.** This figure illustrates the life cycle dynamics of the LP's portfolio after conditioning on business cycle conditions. Panels A through D plot the outcomes for newcommitments, uncalled commitments, liquid wealth, and the stock allocation, respectively, where all outcomes are shown relative to total wealth. The solid (dashed) line shows the average outcome conditional on a expansionary (recessionary) state. Shaded regions indicated 95% confidence intervals, conditional on the macroeconomic state.

The most notable difference across the business cycle is the difference in the LP's stock allocation, particularly during the ramping-up phase (see panel D of Figure 15). For instance, in year 1, stock allocations average 62% of total wealth during expansions compared to 38% during recessions. This difference narrows to about 2% of total wealth by the maintenance phase from year 5 onwards.

There is a smaller difference in new commitments across the business cycle (see panel A). For example, During the ramping-up phase in the first three years, new commitments are lower during recessions compared to expansions by about 5% of total wealth on average. Similarly, panels B and C show that the difference in uncalled commitments and liquid wealth over the business cycle is smaller on average compared to the differences in stock allocation. For example, the difference in PE allocation across the business cycle averages 2% of total wealth over the lifetime of the fund.

The main reason for the pronounced difference in stock allocation across the business cycle, as opposed to PE allocation, is due to adjustment costs. PE allocation involves significant

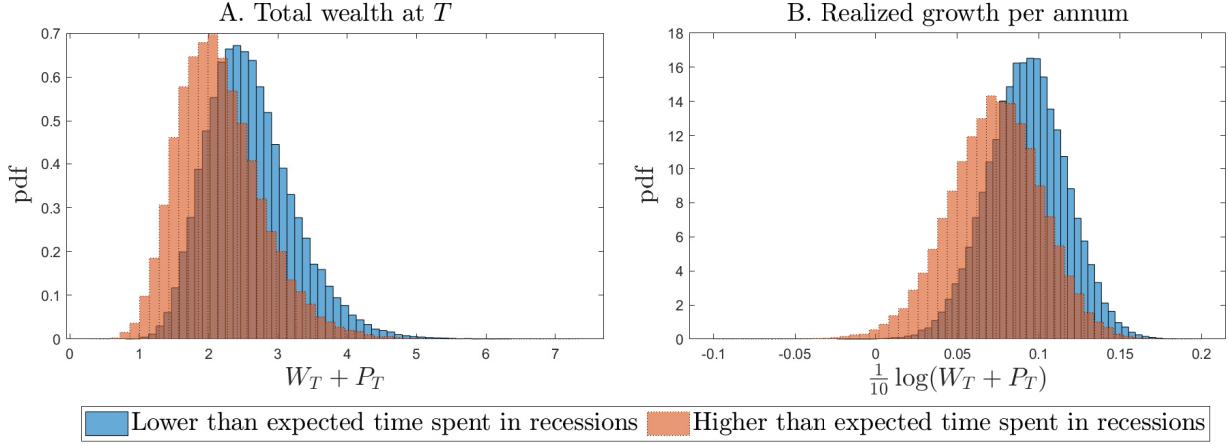


Figure 16: **Business cycles and total wealth outcomes.** This figure illustrates the distribution of total wealth outcomes at the end of the investment horizon  $t = T$ . The initial total wealth is  $W_0 + P_0 = 1$ . Panel A plots the distribution for total wealth  $W_T + P_T$ . Panel B plots the distribution for the realized growth in total wealth per annum,  $\log(W_T + P_T)/10$ . The solid (dashed) bars plot the outcomes conditional on the LP experiencing a lower (higher) than expected number of recessions over its investment horizon.

adjustment costs, whereas adjustment costs for stocks are negligible. Therefore, the LP keeps the PE allocation relatively stable across the business cycle and primarily adjusts its exposure to aggregate risks through stock allocation.

Figure 16 plots the distribution of outcomes at the terminal date  $T$  conditional on business cycle conditions encountered over the LP's investment horizon. The transition probabilities for the macroeconomic state from Table 1 imply that recessions occur  $1/6$  of the time. The solid blue (dashed red) bars plot outcome distributions conditional on the LP encountering recessions for less (more) than  $1/6$  of the time over its investment horizon. Panel A plots the outcomes for wealth. On average, the LP ends up with a terminal wealth of 2.60 and 2.17 conditional on experiencing a lower- and higher-than expected number of recessions, respectively. The corresponding standard deviations of terminal wealth are similar: 0.63 and 0.61 conditional on experiencing a lower- and higher-than expected number of recessions, respectively. The outcomes in terms of annualized realized total returns are shown in panel B. Conditional on experiencing a lower than expected number of recessions, the realized annualized returns average 9.29% with a standard deviation of 2.40%. The annualized return averages 7.35% with a standard deviation of 2.83% if a higher than expected number of recessions is encountered instead.

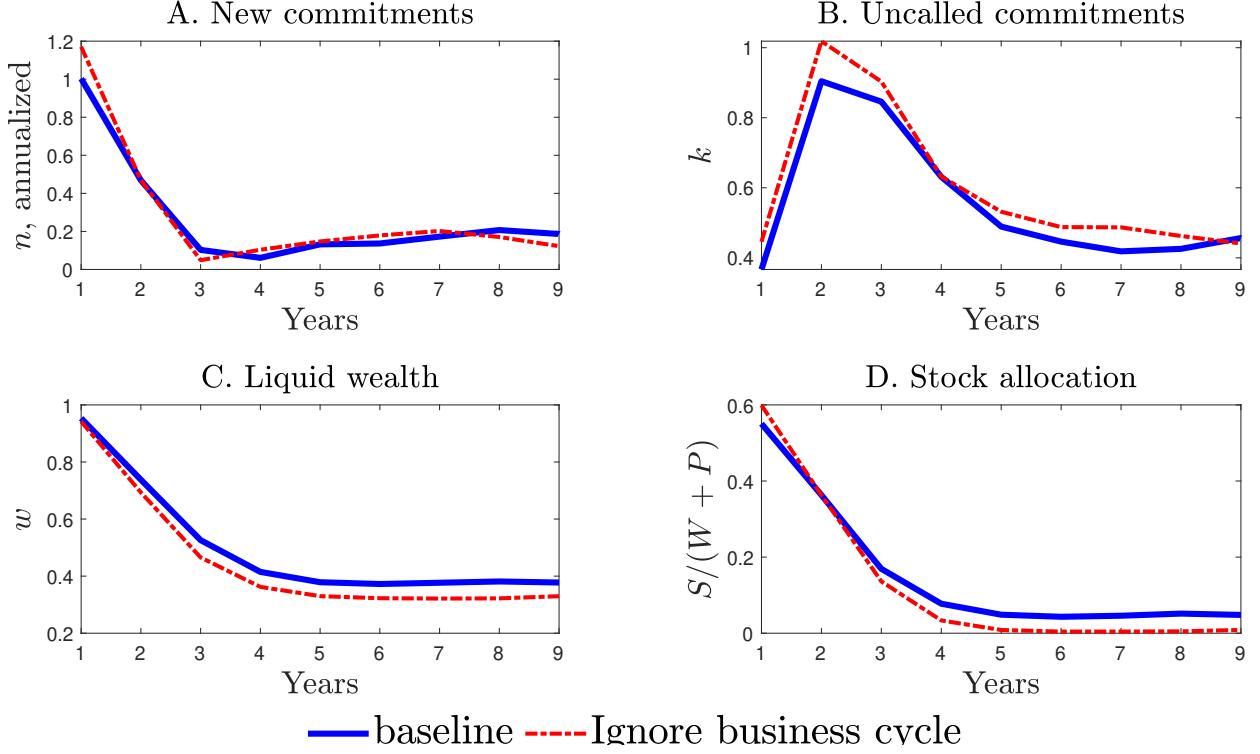


Figure 17: **Life cycle outcomes without accounting for business cycles.** The solid and dashed lines plot the average outcomes when business cycles are accounted for and not accounted for, respectively.

#### 4.5.1 Cost of ignoring business cycle when making PE allocations

Next, we illustrate the importance of accounting for business cycle conditions through the following exercise. We consider a naive LP whose decisions are based on a model that ignores business cycles. Specifically, this naive LP solves the model under the assumption that the economy is always in an expansionary state (i.e., this LP sets  $s_0 = 2$ ,  $p_{21} = 0$ , and uses the remaining parameters from Table 1). The naive LP then applies the resulting policy functions, which ignore business cycle fluctuations, in a setting where such fluctuations are present. We compare the life cycle dynamics of this naive LP to the baseline case, where the LP optimally accounts for business cycle conditions. The results are shown in Figure 17.

We observe that without considering the possibility of recessions, the LP's PE allocation is significantly more aggressive compared to the optimal policy that accounts for business cycle fluctuations. For instance, in year 1, new commitments average 117% of total wealth compared to 100% in the baseline case. During the maintenance phase, the PE allocation averages 67% of total wealth, as opposed to 62% in the baseline case. Consequently, the naive LP defaults substantially more frequently—the naive LP's 10-year cumulative default

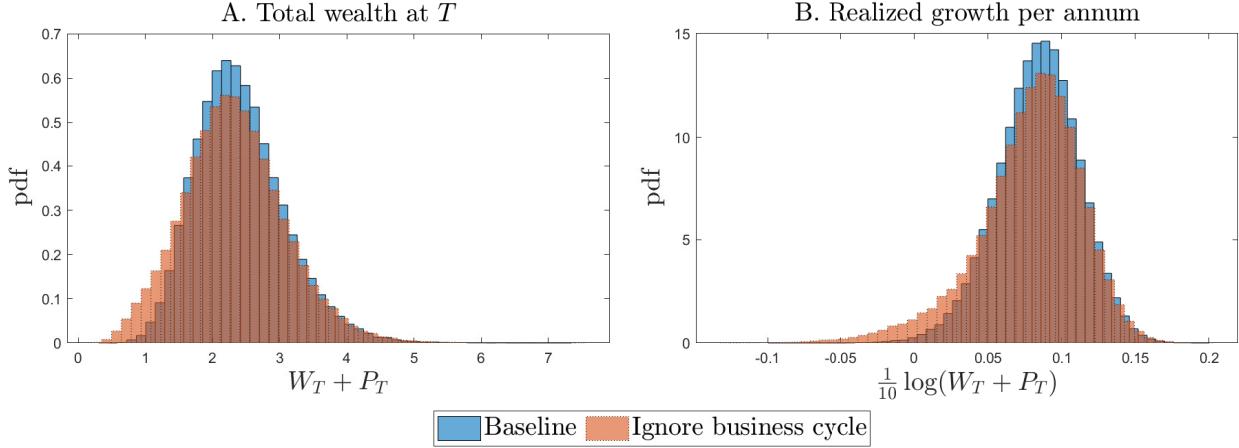


Figure 18: **Total wealth outcome without accounting for business cycles.** This figure illustrates the distribution of total wealth outcomes at the end of the investment horizon  $t = T$ . The initial total wealth is  $W_0 + P_0 = 1$ . Panel A plots the distribution for total wealth  $W_T + P_T$ . Panel B plots the distribution for the realized growth in total wealth per annum,  $\log(W_T + P_T)/10$ . The solid and dashed bars plot the outcomes for the baseline and a naive investor that ignores business cycles, respectively.

probability is 13.6% compared to only 0.1% under the optimal policy.

This large increase in default frequency has a significant negative impact on the naive LP's final portfolio outcomes. Figure 18 compares the outcome distributions for the naive investor against the baseline; panels A and B compare outcomes for terminal total wealth and annualized realized returns, respectively. Compared to the baseline investor that takes business cycles into account, the wealth total wealth distribution for the naive LP has a lower mean (2.32 vs. 2.41) and a higher volatility (0.75 vs 0.66). Similarly, the naive LP's annualized realized returns have a lower mean (7.83% vs 8.41%) and a higher volatility (3.58% vs 2.78%). The biggest difference in outcomes, however, is heavier left-tail encountered by the naive investor. For example, the first and fifth percentiles for the naive LP's realized annualized returns equal -3.06% and 1.14%, as opposed to 1.34% and 3.68% for the baseline case. The naive LP's suboptimal policies translate into a 9.3% loss in initial total wealth—the naive LP's certainty equivalent value is 2.04 compared to 2.25 for the baseline case.

These results underscore the importance of considering business cycle conditions when making optimal PE allocation decisions.

## 4.6 Does serial correlation in PE returns matter for *long-term* investors?

Positive serial correlation is a hallmark feature of the observed returns of illiquid alternative asset, PE investments included. A large literature investigates the sources of this serial correlation and its implications for investors (see, e.g., [Geltner 1991](#) and [Getmansky, Lo, and Makarov 2004](#)). A likely explanation is that serial correlation can arise from having to mark-to-market illiquid assets, which can lead to observed returns that appear smooth even when the underlying true returns do not have serial correlation ([Getmansky, Lo, and Makarov, 2004](#)). As a result, the literature has argued for “unsmoothing” such returns to obtain a more accurate picture of the underlying asset’s risk and return characteristics; for example, naively using smoothed returns can lead to overestimation of Sharpe ratios and faulty portfolio allocations.

In this section, we revisit the question of whether unsmoothing is necessary for optimal portfolio choice. We differ from the existing literature in that we take the perspective of a *long-term* investor. Our headline result is that for long-term investors, whether or not returns must be first unsmoothed may be a moot point. The intuition is as follows. First, long-term investors care about properties of long-horizon returns, which are much less affected by short-term mark-to-market fluctuations. Second, even if serial correlation is a true feature of short run PE returns, it may not be exploitable by investors after taking realistic implementation lags and adjustment costs into account. We provide details below.

Consider the following thought experiment. Suppose, for argument’s sake, that the quarterly autocorrelation of PE expected returns of  $\varrho_{P,1} + \varrho_{P,2} = 0.20$  in our baseline calibration is entirely due to smoothing. The LP may then want to base its portfolio choice on unsmoothed returns by considering a model without serial correlation:  $\varrho_{P,1} = \varrho_{P,2} = 0$ . To keep the overall properties of PE returns unchanged, this LP simultaneously recalibrates the other parameters of the PE returns process so that the expected PE return  $\mathbb{E}[\log R_{P,t+1}|s_t]$  and volatility  $\sigma(\log R_{P,t+1}|s_t)$  remain identical to that of the baseline calibration in both of the macroeconomic states. This involves setting  $\nu_P(1) = 0.0051$ ,  $\nu_P(2) = 0.0392$ ,  $\sigma_P(1) = 0.0772$ , and  $\sigma_P(2) = 0.0427$  while keeping the remaining parameters unchanged from their baseline values in Table 1.

Figure 19 compares the resulting difference in life cycle dynamics between the baseline case and the case where the LP ignores serial correlation in PE returns. We see that the LP’s portfolio allocations appear near-identical across the two cases. Furthermore, the two outcomes remain near-identical when we further condition outcomes on the macroeconomic

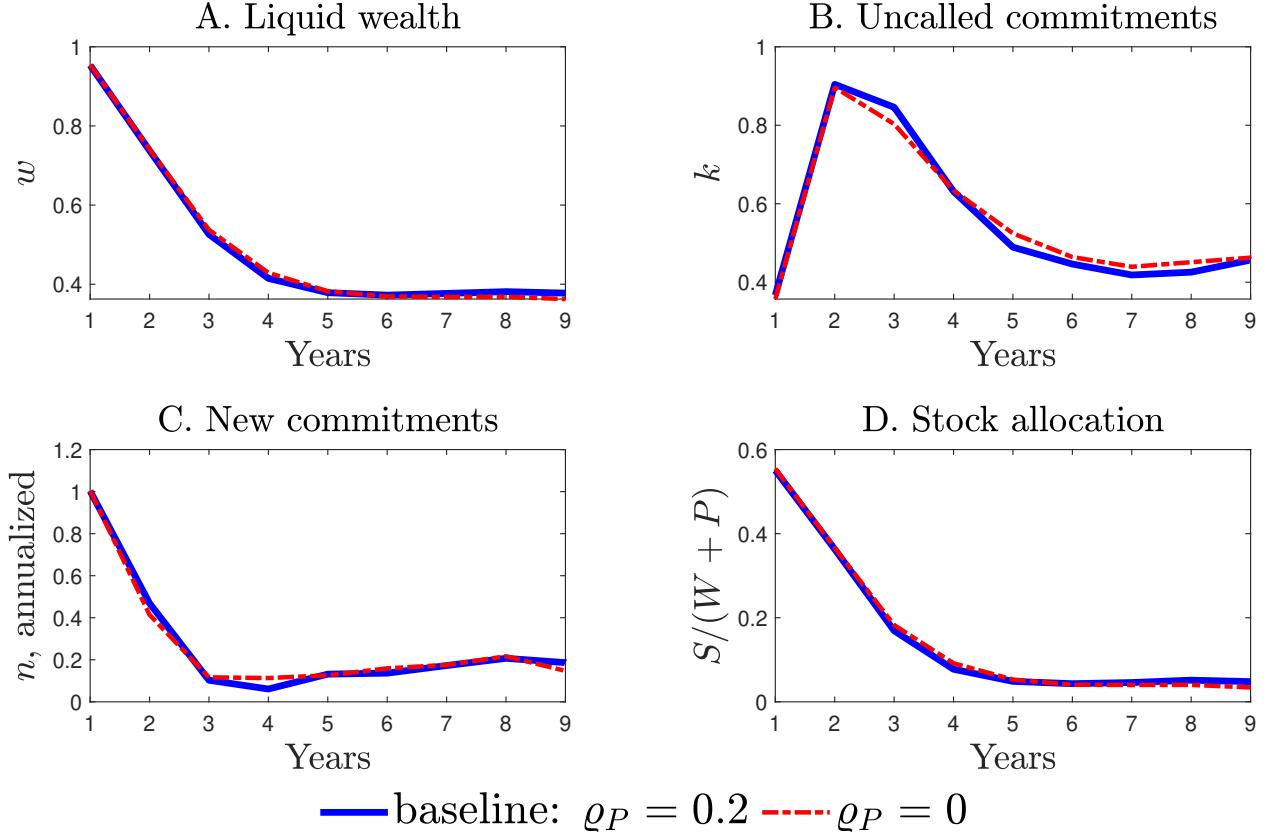


Figure 19: **PE return autocorrelation and life cycle dynamics.** The solid and dashed lines plot the average lifecycle outcomes under (1) our baseline model, and (2) a model where the PE returns process has zero autocorrelation but otherwise have the same return moments.

state (see Figure D.1 in Appendix D).

The results in this section show that for long-term investors, the presence of serial correlation in PE returns may not be a critical factor in determining optimal portfolio allocations.

## 4.7 Risk budget

The risk weights in our baseline calibration  $\theta_S = \theta_P = 1.5$  correspond to 50% risk charges for PE and stocks (the risk budget threshold  $\bar{\theta}$  is normalized to 1). In this section, we conduct comparative static exercises to investigate how the LP's outcomes depend on these risk charges. This is a useful exercise in at least two contexts. First, while risk charges of 50% are representative numbers on average, risk charges can vary widely depending on asset quality. For example, S&P Global assesses anywhere between 35% to 99% market risk charges for equities depending on asset quality (S&P Global, 2023, Table 14). Second, individual LPs may assess different risk charges based on their own individual circumstances. For example,

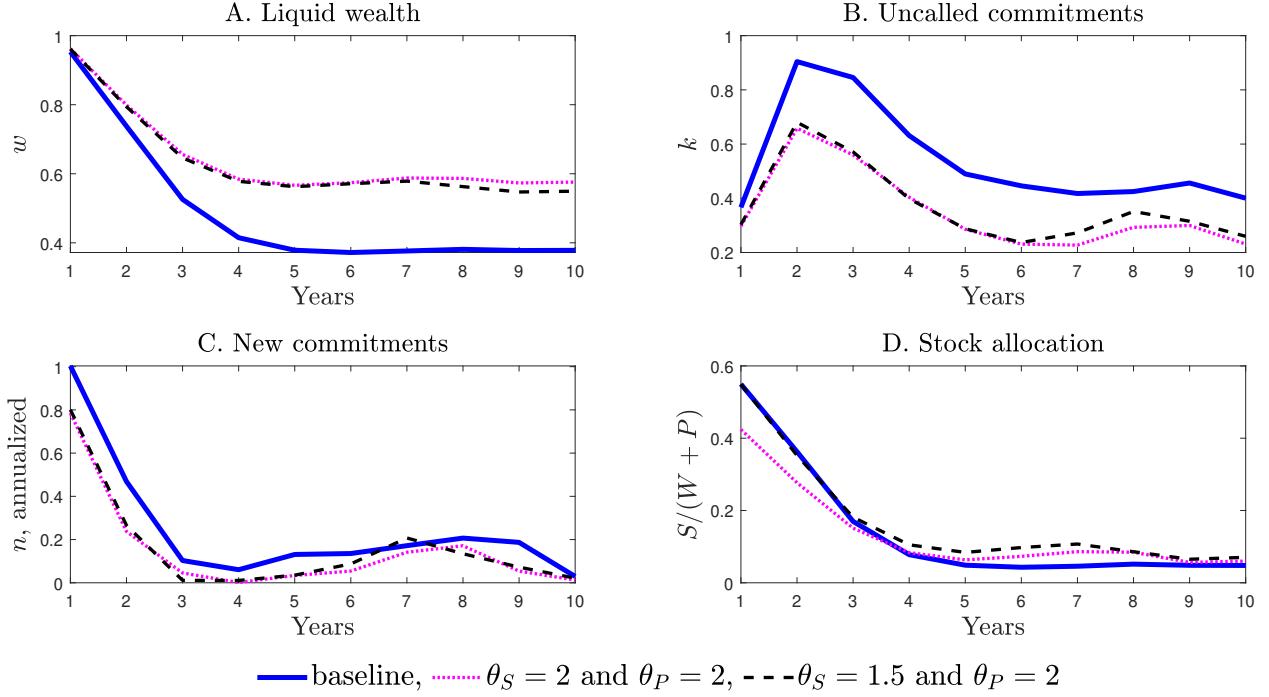


Figure 20: **Life cycle dynamics under different risk charges.** This figure displays average life cycle outcomes under alternative risk charges. For reference, the solid line plots outcomes under the baseline calibration in which  $\theta_P = \theta_S = 1.5$ . The dotted line plots outcomes when both  $\theta_S$  and  $\theta_P$  equal 2, and the dashed line plots outcomes when only  $\theta_P$  is increased to 2.

an insurer may assess higher risk charges for its PE investments depending on properties of its insurance portfolio; the risk weights  $\theta_P$  and  $\theta_S$  can then be interpreted as effective risk charges whose values depend on the riskiness of other components of the insurer’s balance sheet.

We demonstrate the effects of risk charges through two scenarios. In the first scenario, we set  $\theta_S = \theta_P = 2$ , so that risk charges are uniformly higher. In the second scenario, we increase the risk charge for PE to  $\theta_P = 2$  while keeping  $\theta_S = 1.5$  unchanged. For both scenarios, we take the remaining parameters from the baseline calibration in Table 1.

Figure 20 illustrates the impact of risk charges on average life cycle outcomes. Compared to the baseline outcome (solid line), we see that overall allocations to PE and stocks are lower under the first scenario (dotted line) in which the risk charges for both PE and stocks are increased to  $\theta_P = \theta_S = 2$ . For instance, during the initial ramp-up phase, uncalled commitments peak at 66% of total wealth in year 2, as opposed to 90% in the baseline calibration. Similarly, stock allocation in year 1 is reduced to 42% of total wealth on average compared to 55% in the baseline. In the maintenance phase from year 5 onwards, the PE allocation is 42% of total wealth, down from 62% in the baseline. The LP does, however,

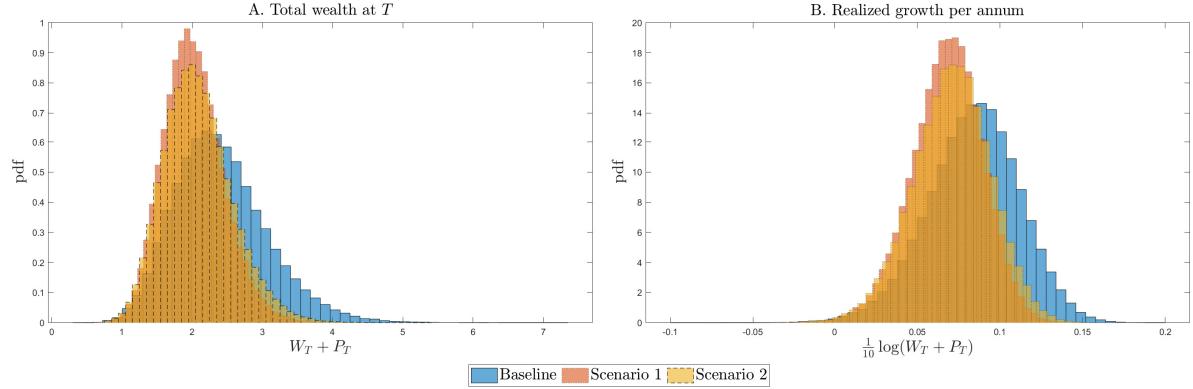


Figure 21: **Total wealth outcomes under different risk charges.** This figure illustrates the distribution of total wealth outcomes at the end of the investment horizon  $t = T$ . The initial total wealth is  $W_0 + P_0 = 1$ . Panel A plots the distribution for total wealth  $W_T + P_T$ . Panel B plots the distribution for the realized growth in total wealth per annum,  $\log(W_T + P_T)/10$ . Both panels compare outcomes under (1) the baseline in which  $\theta_P = \theta_S = 1.5$ , (2) scenario 1 in which both  $\theta_P$  and  $\theta_S$  are increased to 2, and (3) scenario 2 in which only  $\theta_P$  is increased to 2.

compensates with a slightly higher stock allocation during the maintenance phase: 7% of total wealth on average compared to 5% on average for the baseline. Figure D.2 in Section D additionally displays outcomes for the first scenario after further conditioning on the business cycle.

The second scenario isolates the effect of risk charges by only increasing  $\theta_P$ . In this case, we see from Figure 20 that only PE allocations are affected to first order. Specifically, the average outcome for stock allocation remains roughly unchanged from the baseline (compare the solid and dashed lines in panel D) while the outcomes for PE allocation are similar to that under the first scenario (compare the dotted and dashed lines in panels A, B, and C). Figure D.3 in Section D additionally displays outcomes for the second scenario after further conditioning on the business cycle.

Figure 21 compares the distribution of final outcomes under the different risk charge scenarios. We see that higher risk charges do indeed decrease the riskiness of the LP's overall portfolio. For example, the standard deviation of annualized realized returns decreases to 2.35% (from the baseline value of 2.78%) under the second scenario when  $\theta_P$  increases to 2; it further decreases to 2.14% under the first scenario when both  $\theta_P$  and  $\theta_S$  increase to 2. This decrease in risk does come at the expense of lower returns. The average annualized return decreases to 7.08% under the second scenario (down from 8.41% for the baseline), and further decreases to 6.79% under the first scenario.

The costs of imposing higher risk charges in certainty equivalent terms are as follows. Increasing  $\theta_P$  alone to 2 (i.e., scenario 2) reduces the certainty equivalent wealth from the

baseline value of 2.25 to 2.00 or, equivalent, a 11% reduction in certainty equivalent wealth. When we additionally increase  $\theta_S$  to 2 (i.e., scenario 1), the certainty equivalent wealth further decreases to 1.92 for a 15% reduction in certainty equivalent wealth compared to the baseline.

## 5 Conclusion

Exciting recent developments in machine learning have enabled the solution of portfolio choice problems featuring higher dimensionality, more realistic trading frictions, and more complex constraints. In this paper, we employ Deep Kernel Gaussian Processes to accurately characterize the optimal policies in a private asset allocation model that incorporates a host of real-world complexities, including illiquidity, commitment lags, serial correlation in returns, business cycle fluctuations, and regulation-induced constraints. Our approach not only provides novel insights into private equity investing but also offers a flexible blueprint for tackling similarly challenging problems in economics and finance that involve substantial nonlinearities and a large number of state variables.

## References

Ang, A., D. Papanikolaou, and M. Westerfield, 2014, “Portfolio Choice with Illiquid Assets,” *Management Science*, 60, 2737–2761.

Bolton, P., H. Chen, and N. Wang, 2011, “A Unified Theory of Tobin’s  $q$ , Corporate Investment, Financing, and Risk Management,” *Journal of Finance*, 66, 1545–1578.

Bolton, P., H. Chen, and N. Wang, 2013, “Market Timing, Investment, and Risk Management,” *Journal of Financial Economics*, 109, 40–62.

Dimmock, S. G., N. Wang, and J. Yang, 2024, “The Endowment Model and Modern Portfolio Theory,” *Management Science*, 70, 1554–1579.

Duarte, V., D. Duarte, and D. H. Silva, 2024, “Machine Learning for Continuous-Time Finance,” *Review of Financial Studies*, 37, 3217–3271.

Duarte, V., J. Fonseca, A. S. Goodman, and J. A. Parker, 2022, “Simple Allocation Rules and Optimal Portfolio Choice Over the Lifecycle,” Working Paper 29559, National Bureau of Economic Research.

Eling, M., and I. Holzmuller, 2008, “An Overview and Comparison of Risk-Based Capital Standards,” *Journal of Insurance Regulation*, 26.

Gaegau, L., S. Scheidegger, and F. Trojani, 2023, “A comprehensive machine learning framework for dynamic portfolio choice with transaction costs,” Working Paper.

Geltner, D. M., 1991, “Smoothing in appraisal-based returns,” *Journal of Real Estate Finance and Economics*, 4, 327–345.

Getmansky, M., A. W. Lo, and I. Makarov, 2004, “An econometric model of serial correlation and illiquidity in hedge fund returns,” *Journal of Financial Economics*, pp. 529–609.

Giommetti, N., and M. Sorensen, 2024, “Optimal Allocation to Private Equity,” Working Paper.

Gourier, E., L. Phalippou, and M. Westerfield, 2024, “Capital Commitment,” *Journal of Finance*, 79, 3407–3457.

Kingma, D. P., and J. Ba, 2015, “Adam: A Method for Stochastic Optimization,” *arXiv preprint arXiv:1412.6980*.

Korteweg, A., and M. M. Westerfield, 2022, “Asset Allocation with Private Equity,” *Foundations and Trends(R) in Finance*, 13, 95–204.

McKinsey & Company, 2024, “McKinsey Global Private Markets Review 2024: Private markets: A slower era,” Discussion paper, McKinsey & Company.

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, 158–181.

Neuberger Berman, 2022, “The Historical Impact of Economic Downturns on Private Equity,” Discussion paper, Neuberger Berman.

Rasmussen, C. E., and C. K. I. Williams, 2005, *Gaussian Processes for Machine Learning*, The MIT Press.

Renner, P., and S. Scheidegger, 2018, “Machine Learning for Dynamic Incentive Problems,” Working Paper.

Scheidegger, S., and I. Bilionis, 2019, “Machine learning for high-dimensional dynamic stochastic economies,” *Journal of Computational Science*, 33, 68 – 82.

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

S&P Global, 2023, “Insurer Risk-Based Capital Adequacy Methodology And Assumptions,” <https://www.spglobal.com/ratings/en/research/articles/231115-criteria-insurance-general-insurer-risk-based-capital-adequacy-methodology-and-assumptions-12862217>, Accessed: January 2025.

Takahashi, D., and S. Alexander, 2002, “Illiquid Alternative Asset Fund Modeling,” *Journal of Portfolio Management*, 28, 90–100.

Wilson, A. G., Z. Hu, R. Salakhutdinov, and E. P. Xing, 2016, “Deep Kernel Learning,” in Arthur Gretton, and Christian C. Robert (ed.), *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, , vol. 51 of *Proceedings of Machine Learning Research*, pp. 370–378, Cadiz, Spain. PMLR.

# Appendix

## A Specification for expected PE returns

In this section, we use the [Getmansky, Lo, and Makarov \(2004\)](#) model of smoothed returns for illiquid alternative investments to motivate our specification (14) for PE expected returns.

Applied to our context, the [Getmansky, Lo, and Makarov \(2004, equation 21\)](#) model posits that the observed return on PE investments  $\log(R_{P,t+1})$  is a weighted average of contemporaneous and lagged values of the true but unobserved return on PE investments:

$$\log(R_{P,t+1}) = \omega_0 \log(\tilde{R}_{P,t+1}) + \omega_1 \log(\tilde{R}_{P,t}) + \dots + \omega_k \log(\tilde{R}_{P,t+1-k}), \quad (\text{A.1})$$

In terms of notation for this section, we use a tilde to denote quantities associated with the true but unobserved returns. For example,  $\log(R_{P,t+1})$  is the observed PE return at  $t+1$  while  $\log(\tilde{R}_{P,t+1})$  is the unobserved contemporaneous true return. The coefficients  $\omega_j$  in equation (A.1) are positive weights that sum to 1. We use the geometric weighting scheme considered in [Getmansky, Lo, and Makarov \(2004, equation 43\)](#):

$$\omega_0 = (1 - \phi)/(1 - \phi^{k+1}), \quad \omega_j = \omega_0 \phi^j \text{ for } j = 1, \dots, k, \quad (\text{A.2})$$

where  $\phi \in (0, 1)$ . This weighting scheme leads to the convenient AR(1) process for the expected returns on observed PE returns (14).

Specifications (A.1) and (A.2) imply that the expected return for observed PE returns  $\mu_{P,t} \equiv \mathbb{E}_t[\log(R_{P,t+1})]$  evolves as follows:

$$\begin{aligned} \mu_{P,t} &= \omega_0 \mathbb{E}_t[\log(\tilde{R}_{P,t+1})] + \phi \mathbb{E}_t[\omega_0 \log(\tilde{R}_{P,t}) + \dots + \omega_0 \phi^{k-1} \log(\tilde{R}_{P,t+1-k})] \\ &= \omega_0 \mathbb{E}_t[\log(\tilde{R}_{P,t+1})] + \phi \mathbb{E}_t[\log(R_{P,t}) - \omega_0 \phi^k \log(\tilde{R}_{P,t-k})] \\ &= \omega_0 \tilde{\mu}_{P,t} + \phi \log(R_{P,t}) - \omega_0 \phi^{k+1} \mathbb{E}_t[\log(\tilde{R}_{P,t-k})] \end{aligned} \quad (\text{A.3})$$

$$\approx \omega_0 \tilde{\mu}_{P,t} + \phi \log(R_{P,t}) \quad (\text{A.4})$$

$$= \omega_0 \tilde{\mu}_{P,t} + \phi \mu_{P,t-1} + \phi \sigma_{P,t-1} \varepsilon_{P,t}. \quad (\text{A.5})$$

Equation (A.3) is a direct consequence of specifications (A.1) and (A.2);  $\tilde{\mu}_{P,t} \equiv \mathbb{E}_t[\log(\tilde{R}_{P,t+1})]$  is the expected return for the true return process. The approximation (A.4) holds if  $k$  is large enough or if  $\phi$  is small enough. Equation (A.5) writes the observed return as  $\log(R_{P,t}) = \mu_{P,t-1} + \sigma_{P,t-1} \varepsilon_{P,t}$  where  $\mu_{P,t-1} \equiv \mathbb{E}_{t-1}[\log(R_{P,t})]$ ,  $\sigma_{P,t-1} \equiv \sqrt{\text{Var}_{t-1}[\log(R_{P,t})]}$ , and  $\varepsilon_{P,t} \equiv (\log(R_{P,t}) - \mu_{P,t-1})/\sigma_{P,t-1}$ . Comparing equation (A.5) to equation (14), we see that the AR(1) process for expected PE returns (14) is a more flexible version of equation (A.5) in which the loading on the  $\sigma_{P,t-1} \varepsilon_{P,t}$  term is allowed to differ from the autoregressive coefficient.

We make further functional form assumptions in our implementation. Specifically, the specification in equations (12) and (14) assumes (1) the expected true return  $\tilde{\mu}_{P,t} = \tilde{\mu}_P(s_t)$  depends only on the macroeconomic state (hence  $\nu_P(s_t)$  in equation (14) corresponds to the  $\omega_0 \tilde{\mu}_{P,t}$  term from equation (A.5)), (2) the volatility of observed PE returns  $\sigma_{P,t-1} = \sigma_P(s_{t-1})$  depends only on the macroeconomic state, and (3) the shock to observed PE returns  $\varepsilon_{P,t} \sim \mathcal{N}(0, 1)$  is normally distributed.

## B Scaled value functions

**Scaled choice variables.** We work with the following scaled versions of the choice variables:

$$n \equiv \frac{N}{W + P}, \quad (\text{B.1})$$

$$\tilde{\phi}_S \equiv \frac{S + \gamma_S(W + P) \left( \frac{S}{W + P} \right)^2}{W - \gamma_N(W + P)(n - \bar{n})^2}. \quad (\text{B.2})$$

Here,  $n$  is the new PE commitments scaled by total wealth. The interpretation for  $\tilde{\phi}_S$  is as follows. The denominator in equation (B.2) is the amount of liquid wealth that is available to be invested in stocks and bonds after deducting adjustment costs for new PE commitments from current liquid wealth. Of this amount,  $\tilde{\phi}_S$  is the fraction invested in stocks inclusive of stock adjustment costs. We also denote by

$$\phi_S \equiv \frac{S}{W - \gamma_N(W + P)(n - \bar{n})^2}. \quad (\text{B.3})$$

the fraction invested in stocks excluding stock adjustment costs. The relation between  $\tilde{\phi}_S$  and  $\phi_S$  is

$$\phi_S = \frac{-1 + \sqrt{1 + 4\gamma_S[w - \gamma_N(n - \bar{n})^2]\tilde{\phi}_S}}{2\gamma_S[w - \gamma_N(n - \bar{n})^2]}. \quad (\text{B.4})$$

**Scaled value function in default.** The recursive formulation for the scaled value function in default is

$$v_D(t, s) = \max_{\tilde{\phi}_S \in [0, \tilde{\phi}_{S,max}(s)]} \mathbb{E} \left[ (g'_D)^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}} \mathbb{E} \left[ v_D(t+1, s')^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}} \quad (\text{B.5})$$

where

$$g'_D \equiv \frac{W'}{W} = R'_S \phi_S + R_f(s) \left( 1 - \tilde{\phi}_S \right) - \Gamma(\theta_S \tilde{\phi}_S) \quad (\text{B.6})$$

is the growth rate of total wealth in default and

$$v_D(T, s) = 1$$

is the terminal condition. Note that when computing the scaled value function (B.5), equations (B.2), (B.3), and (B.4) for  $\tilde{\phi}_S$ ,  $\phi_S$ , and their relation, respectively, are computed with  $P = 0$  and  $n = 0$ . In addition, note that while the  $\mathbb{E} [u_D(t+1, s')^{1-\gamma} |s|]^{\frac{1}{1-\gamma}}$  on the right-hand side of equation (B.5) affects the scaled value function, it does not affect the optimal portfolio choice. This is because the growth rate of total wealth (B.6) does not depend on  $s'$ . As a result, the optimal policies in default only depend on the macroeconomic state  $s$ .

The upper bound  $\tilde{\phi}_{S,max}(s) \equiv \sup \left\{ \tilde{\phi}_S \in [0, 1] : \inf_{R'_S} g'_D = R_f(s) \left( 1 - \tilde{\phi}_S \right) - \Gamma(\theta_S \tilde{\phi}_S) \right\}$  appearing in problem (B.5) ensures that the growth in wealth  $g'_D$  is strictly positive so that utilities are well-defined. The upper bound equals

$$\tilde{\phi}_{S,max}(s) = \begin{cases} 1 & \text{if } \theta_S \leq \bar{\theta}, \\ \min \left\{ \frac{\bar{\theta}}{\theta_S} + \frac{-R_f(s) + \sqrt{R_f(s)^2 + 4\kappa\theta_S R_f(s)(\theta_S - \bar{\theta})}}{2\kappa\theta_S^2}, 1 \right\} & \text{if } \theta_S > \bar{\theta}. \end{cases} \quad (\text{B.7})$$

Similarly, the lower bound  $\tilde{\phi}_S \geq 0$  rules out shorting which is also necessary to ensure that utilities are well-defined.

**Scaled value function before default.** The recursive formulation for the scaled value function before default is

$$v(t, w, k, \mu_P, s) = \max_{n, \tilde{\phi}_S} \mathbb{E} \left[ \max \left\{ g'_D v_D(t+1, s'), g' v(t+1, w', k', \mu'_P, s') \right\}^{1-\gamma} | w, k, \mu_P, s \right]^{\frac{1}{1-\gamma}} \quad (\text{B.8})$$

subject to

$$g'_D = (1-w) [\lambda_D(s') + \alpha(s')(1-\lambda_D(s'))] R'_P - \Gamma(\theta_D) \\ + [w - \gamma_N(n - \bar{n})^2] [\phi_S R'_S + (1 - \tilde{\phi}_S) R_f(s)],$$

$$\theta_D = \frac{\theta_S [w - \gamma_N(n - \bar{n})^2] \tilde{\phi}_S}{1 - \gamma_N(n - \bar{n})^2},$$

$$g' = (1-w) R'_P + [w - \gamma_N(n - \bar{n})^2] [\phi_S R'_S + (1 - \tilde{\phi}_S) R_f(s)] - \Gamma(\theta), \quad (\text{B.9})$$

$$\theta = \frac{\theta_P(1-w) + \theta_S [w - \gamma_N(n - \bar{n})^2] \tilde{\phi}_S}{1 - \gamma_N(n - \bar{n})^2},$$

$$k' = \frac{[1 - \lambda_K(s')]k + [1 - \lambda_N(s')]n}{g'},$$

$$w' = \frac{\left\{ \begin{array}{l} \lambda_D(s') R'_P (1-w) - \lambda_K(s') k - \lambda_N(s') n - \Gamma(\theta) \\ + [w - \gamma_N(n - \bar{n})^2] [\phi_S R'_S + (1 - \tilde{\phi}_S) R_f(s)] \end{array} \right\}}{g'},$$

$$\mu'_P = \varrho_{P,1} \mu_P + \varrho_{P,2} \log R'_P + \nu_P(s'),$$

$$n \in [n_{\min}(w), n_{\max}(w)], \quad \tilde{\phi}_S \in [0, \tilde{\phi}_{S,\max}(w, n, s)], \quad (\text{B.10})$$

with terminal condition

$$v(T, w, k, \mu_P, s) = 1.$$

Here,  $g'_D = W'_D/(W + P)$  and  $g' = (W' + P')/(W + P)$  denote the growth rates of total wealth conditional on defaulting and not defaulting, respectively.

The limits on the choice variables (B.10) ensure that utility is well-defined. To see why, note that the log returns on PE and the stock are normally distributed and are therefore unbounded. Hence, this would require  $\max\{g'_D, g'\} > 0$  for all realization of return shocks. This condition implies  $\inf g'_D = [w - \gamma_N(n - \bar{n})^2](1 - \tilde{\phi}_S) R_f(s) - \Gamma(\theta_D) \geq 0$  where the infimum is taken over both  $R'_P$  and  $R'_S$ . From this, we obtain

$$n_{\min}(w) = \max \left\{ \bar{n} - \sqrt{\frac{w}{\gamma_N}}, 0 \right\}, \quad n_{\max}(w) = \bar{n} + \sqrt{\frac{w}{\gamma_N}},$$

and

$$\tilde{\phi}_{S,max}(w, n, s) = \begin{cases} 1 & \text{if } \tilde{\theta}_S(w, n) \leq \bar{\theta}, \\ \min \left\{ \frac{\bar{\theta}}{\tilde{\theta}_S(w, n)} + \frac{-\tilde{R}_f(w, n, s) + \sqrt{\tilde{R}_f(w, n, s)^2 + 4\kappa\tilde{\theta}_S(w, n)\tilde{R}_f(w, n, s)[\tilde{\theta}_S(w, n) - \bar{\theta}]}}{2\kappa\tilde{\theta}_S(w, n)^2}, 1 \right\} & \text{if } \tilde{\theta}_S(w, n) > \bar{\theta}, \end{cases}$$

where

$$\tilde{\theta}_S(w, n) \equiv \frac{\theta_S[w - \gamma_N(n - \bar{n})^2]}{1 - \gamma_N(n - \bar{n})^2}, \quad \text{and } \tilde{R}_f(w, n, s) \equiv [w - \gamma_N(n - \bar{n})^2]R_f(s).$$

**Scaled value function for the heuristic problem (28).** The value function for the heuristic problem (28) can be scaled as  $V_{\text{heuristic}}(t, W_{\text{total}}, s) = [W_{\text{total}}v_{\text{heuristic}}(t, s)]^{1-\gamma}/(1-\gamma)$  where  $v_{\text{heuristic}}$  solves

$$v_{\text{heuristic}}(t, s) = \max_{\phi_S, \phi_P} \mathbb{E} \left[ (g'_{\text{total}})^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}} \mathbb{E} \left[ v_{\text{heuristic}}(t+1, s')^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}} \quad (\text{B.11})$$

subject to

$$\begin{aligned} g'_{\text{total}} &\equiv \frac{W'_{\text{total}}}{W_{\text{total}}} = R'_S \phi_S + R'_P \phi_P + R_f(s) (1 - \phi_S - \phi_P) - \Gamma(\theta_S \phi_S + \theta_P \phi_P), \\ \phi_S, \phi_P &\in [0, 1], \\ \phi_S + \phi_P &= 1 \end{aligned}$$

and the terminal condition  $v_{\text{heuristic}}(T, s) = 1$ .

## C Numerical Implementation

We first solve the scaled problem in default (B.5). In doing so, we note that the the value function can be split as

$$v_D(t, s) = \mathbb{E} \left[ (u_D(t+1, s'))^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}} \max_{\tilde{\phi}_S \in [0, \tilde{\phi}_{S,max}(s)]} \mathbb{E} \left[ (g'_D)^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}}. \quad (\text{C.1})$$

This is because  $g'_D$  does not depend on  $s'$  (see equation (B.6)). Hence, the optimal portfolio choice problem depends  $s$  alone and is obtained by solving

$$\max_{\tilde{\phi}_S \in [0, \tilde{\phi}_{S,max}(s)]} \mathbb{E} \left[ (g'_D)^{1-\gamma} |s| \right]^{\frac{1}{1-\gamma}}. \quad (\text{C.2})$$

To ensure numerical accuracy, we use adaptive quadrature in to compute expectations.

After obtaining  $v_D(t, s)$ , we solve the scaled problem before default (see equation (B.8)) using backward induction. The procedure is as follows:

1. Initiate the value function at  $t = T$  using the terminal condition  $v(T, w, k, \mu_P, s) = 1$ .

2. Obtain sample points for the state variables  $(w, k, \mu_P)$ . In doing so, we sample 800 points via a Halton sequence over the hypercube  $\{(w, k, \mu_P) : w \in [0, 1], k \in [0, 1.5], \mu_P \in [\mu_{P,lb}, \mu_{P,ub}]\}$  where the bounds  $\mu_{P,lb}$  and  $\mu_{P,ub}$  are set to be the 0.1st and 99.9th percentiles of the stationary distribution for  $\mu_P$ , respectively.
3. Given a value function  $v_{GP}(t + 1, w, k, \mu_P, s)$  that has been previously fitted via Deep Kernel Gaussian Process Regression, do the following:
  - (a) For each  $s$  at time  $t$ , solve problem (B.8) for the sample points for  $(w, k, \mu_P)$  obtained in step 2. When doing so, we use the fitted value function  $v_{GP}(t + 1, \cdot)$  as the continuation value when not defaulting next period. As in the problem after default, we use adaptive quadrature to compute expectations to ensure numerical accuracy. To ensure that we find a global maximum, we first consider a coarse grid of policy points. We then run an interior-point optimization algorithm starting from the best policy based on the coarse grid. The end result is a set of observations for the value and policy functions at the sample points.
  - (b) We use the observations from step (a) to fit deep kernel Gaussian processes for the value and policy functions at time  $t$  for each state  $s$ . Our choice for the deep kernel (24) is as follows. The neural network component  $NN$  takes the three-dimensional input  $(w, k, \mu_P)$  and passes it through 3 hidden layers with 64, 32, and 16 units, respectively, before outputting 2 features (the shape of the hidden layers follow a so-called “decoder structure”); we use Gaussian error linear unit (GELU) activation. These 2 output features are then passed through a Matern 5/2 kernel (20). We train the parameters of the deep kernel Gaussian processes by maximizing the marginal likelihood (23) using the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.001. To ensure that the training converges to a global optimum, we run the Adam optimizer 10 times from different initial points. We always use the trained parameters from step  $t + 1$  as one of the initial points; the remaining initial points are randomly drawn with Xavier initialization for the neural network weights. We standardize both input and output variables when training.
4. Repeat step 3 until  $t = 0$ .

Figure 6b and Figure C.1 illustrate the value function and policy functions at  $t = 0$ , respectively.

## D Additional Figures

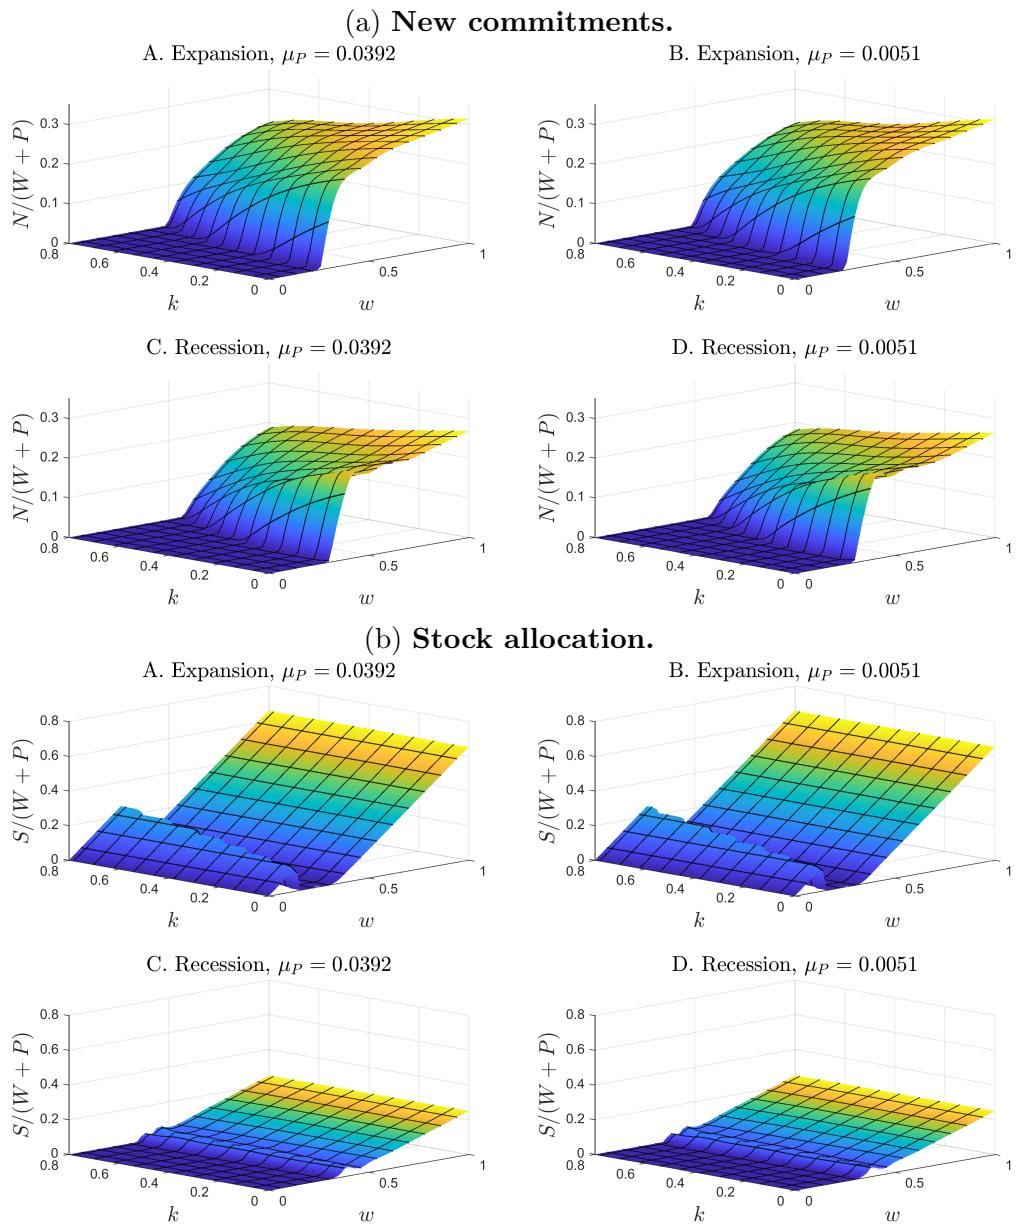


Figure C.1: Policies at  $t = 0$ .

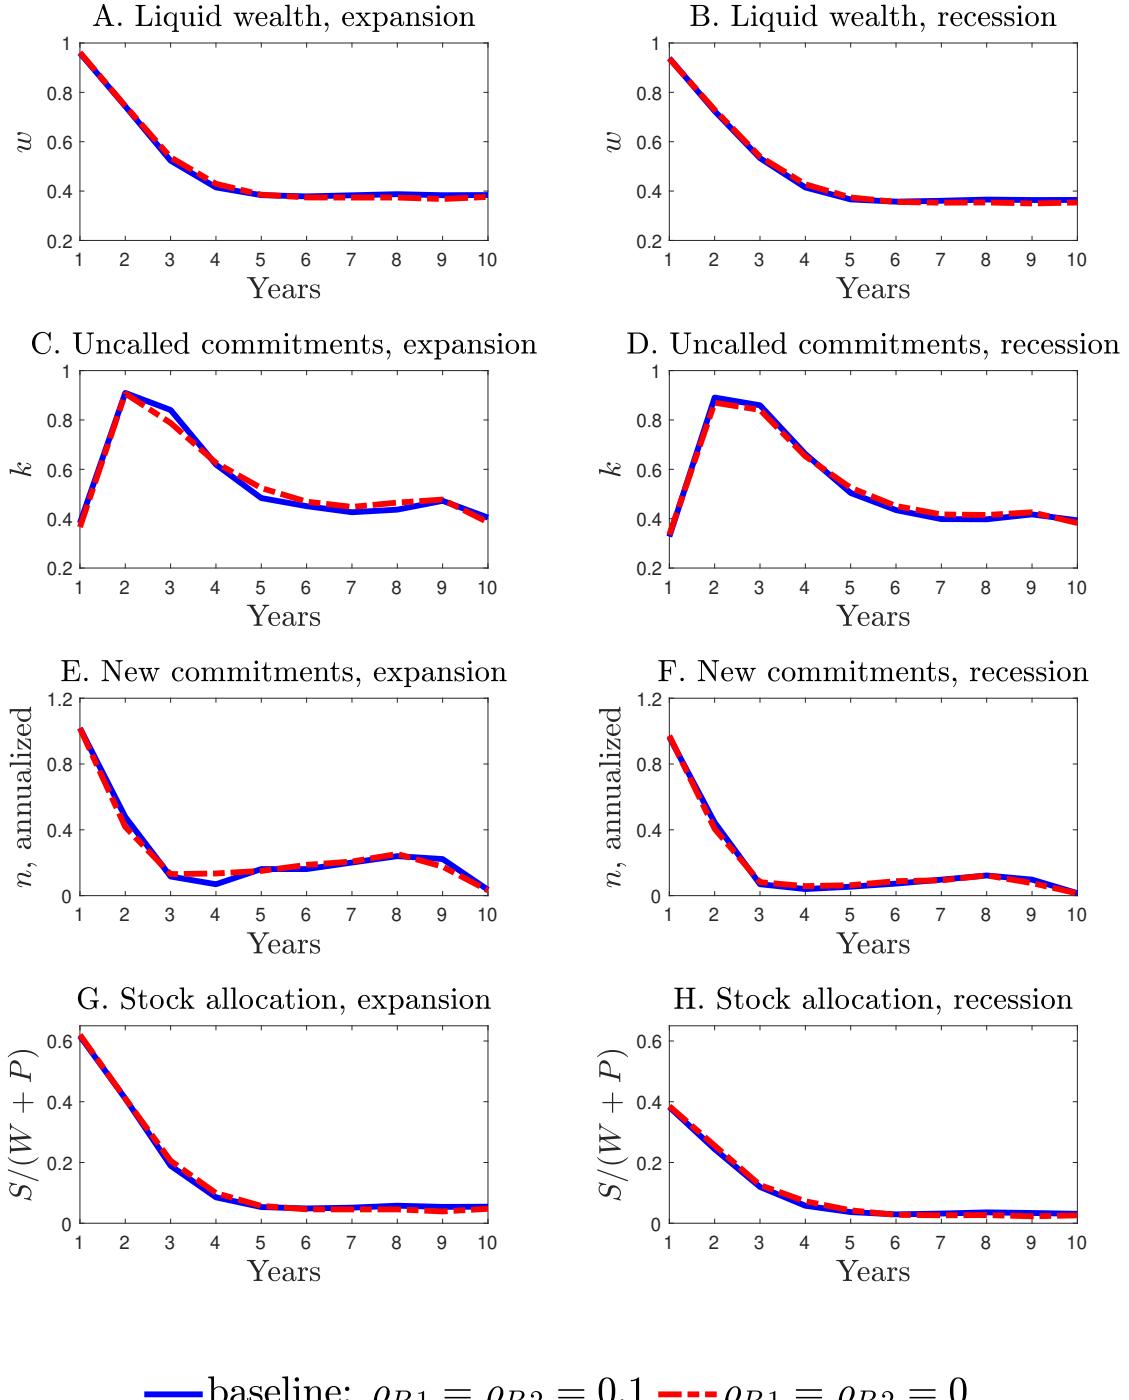


Figure D.1: **Life cycle comparison (conditional), PE return autocorrelation.** The solid and dashed lines plot the average lifecycle outcomes under (1) our baseline model, and (2) a model where the PE returns process has zero autocorrelation but otherwise have the same return moments.

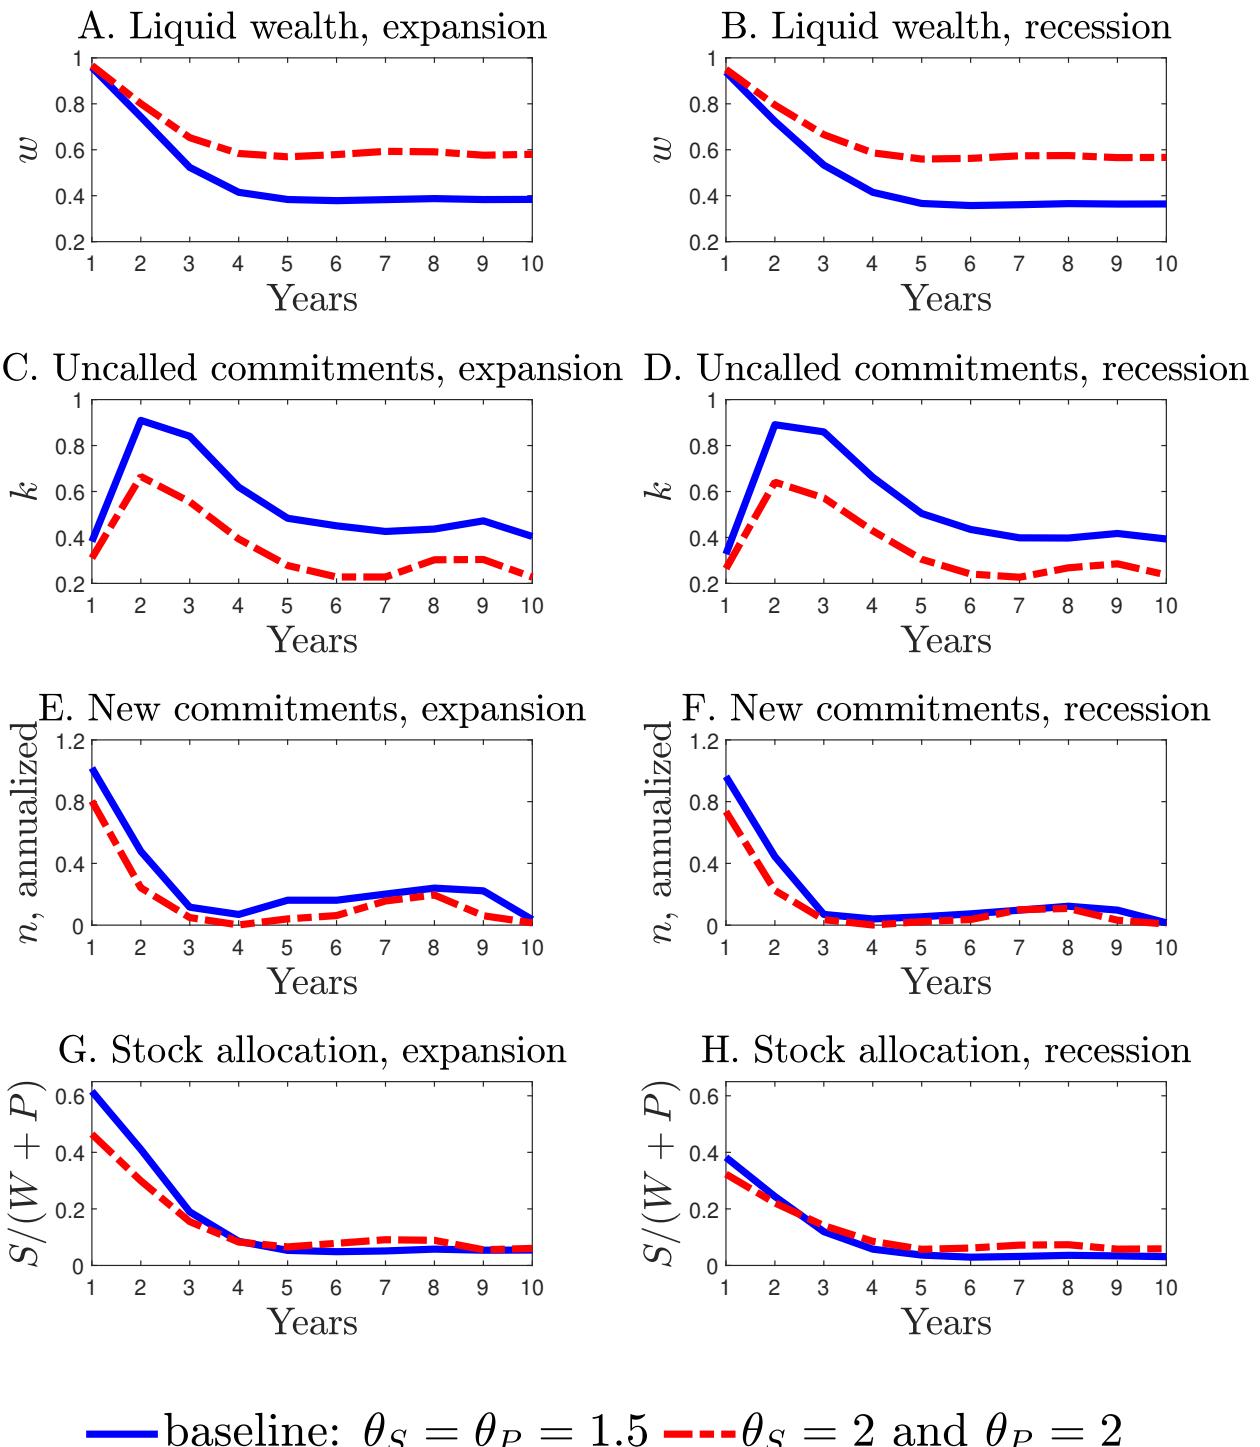


Figure D.2: Life cycle comparison (conditional), risk budget 1.

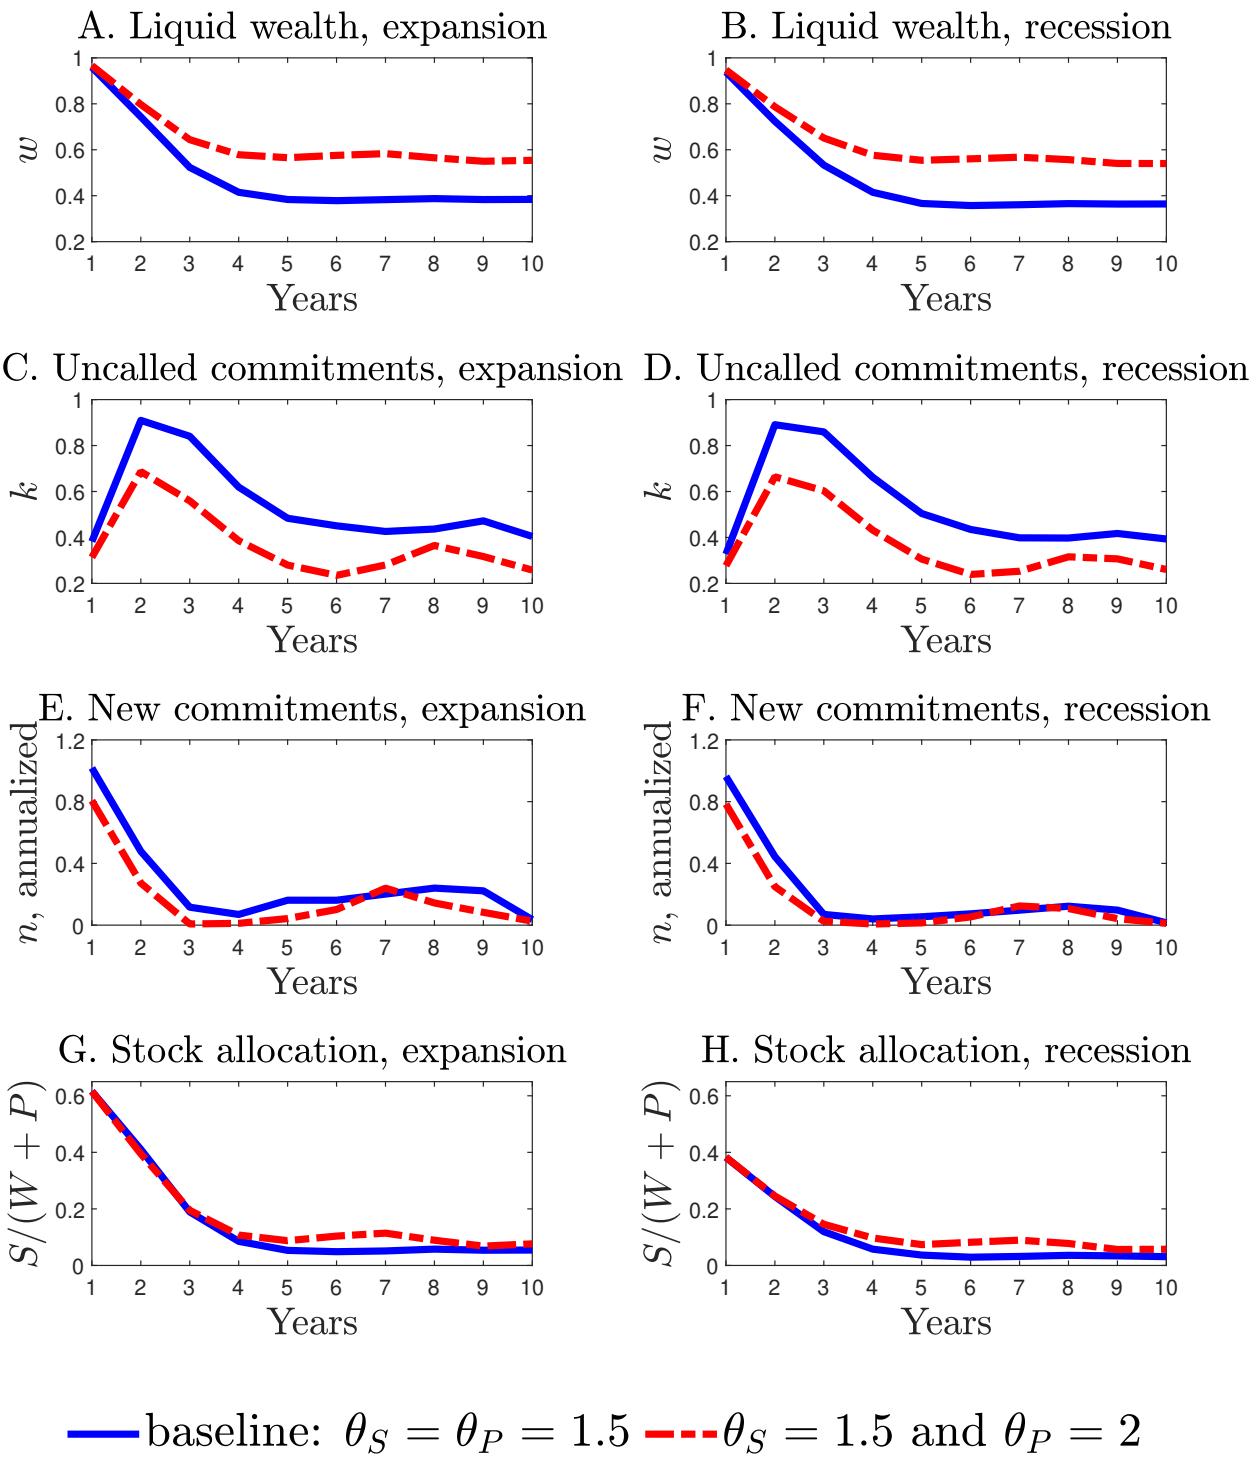


Figure D.3: Life cycle comparison (conditional), risk budget 2.