Strategic Asset Allocation for Sovereign Wealth Funds*

Alexander Michaelides† Yuxin Zhang ‡

March 2022

*We thank Francisco Gomes for detailed feedback and seminar participants at Imperial College London for useful comments. All remaining errors are our own.

†Department of Finance, Imperial College London, South Kensington Campus, London SW7 2AZ, UK. E-mail: a.michaelides@imperial.ac.uk.

‡School of Finance, Renmin University of China, Beijing, 100872, China. E-mail: yx.zhang@ruc.edu.cn.

Strategic Asset Allocation for Sovereign Wealth Funds

Abstract

We propose a normative strategic asset allocation model for long horizon sovereign wealth funds (SWFs). SWFs optimally choose asset allocation to smooth background shocks from persistent commodity prices and business cycles subject to a sustainable spending constraint. SWFs should have lower stock market exposure as they mature, and should use the SWF to smooth business cycle shocks, rather than commodity price shocks, in well-diversified economies. The government's taxation capacity, or consumption commitments, substantially affect asset allocation, and low real interest rates generate a "reaching for yield" portfolio behavior. Higher expected growth without a sustainability constraint generates a smaller SWF.

JEL Classification: E32, G11, G18.

Key Words: Strategic asset allocation, sovereign wealth funds, horizon effects, commodity prices, business cycles.

1 Introduction

How should long term investors like sovereign wealth funds (SWFs) accumulate and decumulate reserves and invest their endowment to achieve their long term and short run objectives? Despite the proliferation of SWFs across countries and across size, the literature providing a normative analysis of SWF management is relatively scant. We think this partly arises because of data limitations emphasized by, for instance, Bernstein et al. (2013), and because the topic shares similarities with the life cycle portfolio choice literature where non-tradeable, undiversifiable labor income can be thought of as the equivalent of the natural resource.[1] The natural resource acts like an implicit riskless asset, just like labor income from an investor's point of view, and the SWF can take more aggressive positions early in its creation, and move into bonds as the resource is depleted.[2]

We argue that substantial and important differences exist between the two frameworks. What are the main differences with life-cycle portfolio choice? The primary difference involves the specification of the SWF's objective function and the constraints a sovereign government faces, which are substantially different from the ones faced by a self-interested individual household. Norway's experience validates this motivation. The stated objectives of the Norwegian SWF emphasize its long-term horizon taking into account the welfare of current and future generations but also the SWF's usefulness as a stabilization tool for the Norwegian economy, both in periods of overheating and in periods when the economy needs fiscal stimulus.[3] These objectives imply both a long run perspective in preserving funds for

[1] For instance, see Viceira (2001) for a long-horizon model and Cocco et al. (2005) for a more explicit life-cycle model with these portfolio implications.

[2] Ang (2014) discusses this point (pp. 167-170).

[3] See https://www.nbim.no/en/the-fund/about-the-fund/.

future generations that are yet unborn, but also a short run perspective to stabilize the economy over the business cycle. Therefore, an explicit solution of a normative SWF strategic asset allocation model in the spirit of Campbell and Viceira (2002) that takes into account these objectives and constraints is warranted; Cochrane (2022) points out, for example, that "a sovereign wealth fund might think ... about which streams of payments can bear risk, and which cannot" (p. 14).

To make progress in building a normative strategic asset allocation model for a SWF, we start by explicitly specifying the objective function. This involves the tradeoff between a long horizon to take into account the interests of future generations, and a shorter run objective that depends on the state of the business cycle to stabilize economic fluctuations in a sustainable way. We explicitly include a sustainability constraint where the SWF withdraws the expected SWF returns (see Campbell and Sigalov (2022)). The sustainability constraint generates a tight link between the endogenous asset allocation decision and the withdrawal rate because the asset allocation decision determines expected returns. The sustainability constraint assumption is one way of capturing the preference for preserving the SWF for future generations and reflects the thinking in a successfully-run SWF like the one in Norway.

Importantly, we also need to take into account certain minimum expenditures that the SWF needs to cover (these commitments could be infrastructure projects or public sector wages). It is important to emphasize that w]e think of the SWF as part of the government and therefore governments with different taxation capacities should take into account that aspect of public finances in determining the withdrawal rate (which in our case translates directly to an asset allocation decision).

Finally, in building the SWF normative model, the risks faced by a long-term investor like a SWF need to explicitly specified. One important risk involves the characteristics of the commodity value process; Deaton and Laroque (1992) early on pointed out that the real value of different commodities typically follows an AR(1) in annual data over the twentieth century. Moreover, the simple AR(1) model can sometimes outperform empirically more sophisticated speculative storage models (Deaton and Laroque (1996)). We therefore introduce a persistent commodity price as one key background risk. A second important background risk is the local business cycle, while we assume no time variations in the investment opportunity set. Background uncertainty therefore works through a rich structure of cross correlations across commodity prices, stock returns and the local business cycle and through persistence in commodity prices and economic fluctuations.

We assume away any political or governance problems because we view the pure economic problem sufficiently complex to address, even after all political and governance issues have been dealt with. We also assume that the SWF invests in an internationally diversified stock market index to avoid any potential problems arising from the "Dutch disease" where a country might lose international competitiveness through real exchange rate appreciation if the proceeds are invested domestically. This approach is consistent with the strategy of Norway's SWF, an internationally recognized leader in SWF management.[4] We therefore also leave for future work the interesting issue of international stock selection in the context of SWF management (see Kotter and Lel (2011)).

There are a number of results that arise from our normative model. First, over the long

[4] Ang et al. (2014) provide a recent excellent review of current and potential future investment strategies of Norway's SWF in a report to the Norwegian ministry of finance.

run, strategic asset allocation involves a derisking approach as the SWF matures (equivalently, as the natural resource approaches complete depletion). Moreover, a higher (lower) expected growth rate in the economy implies a higher (lower) desired withdrawal rate through a permanent income hypothesis logic and that is achieved through a higher (lower) average asset allocation in stocks.

Second, over shorter time periods, the asset allocation (and in this setup the resulting withdrawal rate) depends on the correlation between the GDP innovation and the commodity price innovation. If the economy is well diversified and the correlation is low (as would be the case in a country like Norway where SWF revenues do not determine GDP), then short run stabilization is best achieved by reacting to the GDP innovation (the business cycle) rather than the commodity price. In recessions (expansions), a higher (lower) withdrawal is needed and this is achieved with a higher (lower) average asset allocation in stocks. On the other hand, when the economy is not well-diversified (as would be the case in a country like Saudi Arabia where oil revenues materially impact GDP), then we have the opposite prediction in the following sense. Short run stabilization continues to go the same way conditional on the business cycle shock but the effect is quantitatively very small. Instead, what changes is that now the response to the commodity price innovation is much larger; with the higher correlation between the business cycle and commodity prices, a low commodity price innovation implies a higher withdrawal rate, and a higher corresponding asset allocation in stocks to achieve better smoothing. The broad conclusion is that whether asset allocation responds more aggressively to business cycle or commodity price shocks depends critically on the underlying structure of the economy (the correlation between commodity prices and aggregate output).

Third, a key parameter involves the extent to which the SWF is explicitly affected by the state of the business cycle. We model the ability to rely on some tax revenues for government expenditures by introducing explicitly a small component of the exogenously evolving GDP in the objective function. In good (bad) times, the presence of some component of the good (bad) state GDP in the objective function will reduce (increase) the need to take funds out of the SWF. Equivalently, in good times, government expenditures can be paid more easily with existing taxes and therefore there is less of a need to withdraw from the SWF. Asset allocation becomes a lot more conservative when this parameter is set to zero because the government cannot rely at all on any tax revenues arising outside the SWF. The link between asset allocation and withdrawals through the sustainability constraint therefore generates lower average withdrawals. Our modelling approach allows us to quantify this intuition in a simple way; the model can be adjusted to perform alternative experiments and offer advice to SWFs facing these economic conditions.

Fourth, the sustainable spending constraint generates a clear "reaching for yield" preference where a lower risk free rate makes it harder to hit the spending constraint and therefore a higher exposure in the stock market is taken. Despite this higher exposure, the withdrawals are still lower than in the higher real interest rate environment, leaving the average size of the fund unchanged. The prediction that lower interest rates generate "reaching for yield" behavior is similar to the Campbell and Sigalov (2022) insight. In the current setup, the prediction is shown to be robust in an intertemporal setting for a SWF facing background uncertainty risks with the prediction continuing to hold over the whole life of the SWF.

Fifth, in the presence of fixed consumption commitments, the SWF acts in a more risk averse way. Consumption commitments act like a habit in the same way as in Chetty and Szeidl (2007) and generate a more risk averse behavior than in the absence of these commitments. The lower asset allocation in stocks with commitments implies lower withdrawals; interestingly the average wealth accumulation is the same in both cases. It would be interesting to investigate further how such commitments can arise and evolve and how they can affect asset allocation decisions.

Sixth, we experiment with both different types of constraints (limiting the exposure to the stock market and limiting the withdrawals) under the sustainable spending constraint. The results remain qualitatively unchanged from the baseline model in terms of shapes of average allocations over the life of the SWF. What makes a significant change in predictions is when we allow an optimal withdrawal rate without a sustainable spending constraint. In that instance, average SWF size is lower because withdrawals are larger in the early years of the SWF. The lower initial wealth implies more aggressive asset allocations that allow eventually higher withdrawal rates. The effects are further magnified with a higher expected growth rate; in this case withdrawals are even higher over a longer time period and the SWF size is even smaller, generating even more aggressive asset allocations. These findings are to be expected given the results in Carroll (1997): in a similar model with background risk, wealth accumulation increases over the life cycle when expected income/productivity growth is lower, and the economic intuition is essentially identical. The portfolio choice implications follow from the higher wealth accumulation and are again consistent with the literature on background labor income risk (Cocco et al. (2005)). A lower early financial wealth accumulation means that the implicit riskless asset in the form of natural resources is larger proportionately and hence diversification implies a higher allocation in stocks early on.

Our model is related to the recent contribution by Gilbert and Hrdlicka (2015) in the sense that we model the intertemporal choices of a long-horizon investor (in our case the government, in their case a university endowment). We differ by focussing on the choices of the SWF that takes as given the mandate from the government but is then independent to achieve these tasks, while also introducing the sustainable spending constraint and explicitly setting up the problem so that the resource endowment is eventually depleted, something that is not expected to happen with university donations.

The paper is organized as follows. Section 2 describes Norway's experience as it will be used to make certain assumptions about the SWF setup. Section 3 lays out the normative model and Section 4 describes the calibration. Section 5 presents the results from the baseline model and then performs various comparative statics experiments. Section 6 presents what happens when the sustainable spending constraint is replaced by an optimal withdrawals assumption, and Section 7 concludes.

2 Norway's experience

We aim to build a normative, quantitative, strategic asset allocation model for SWFs. Before proceeding further with specifying the model and proceeding with a calibration, we think it is therefore useful to describe the experience of one of the more successful SWFs for which readily available data exist, namely Norway's SWF.

Figure 1 shows the income from natural resources in Norway as a percent of GDP. Norway's SWF was established in 1996 but it is clear from the diagram that substantial revenues as a percent of GDP arose long before 1996. In the 1970s the income was below 2 percent of GDP but this rapidly rose to be between 4 and 8 percent of GDP between 1980 and 1996. The income rapidly rose in the 2000s where it averaged around 10 percent of GDP. This diagram illustrates that SWFs might get started after a relatively sizeable natural resource receipt to GDP is reached; we therefore start our baseline calibration with a five percent natural resource to GDP endowment.

Figure 2 shows the size of the SWF in Norwegian GDP units between 1996 and 2020. Figure 2 shows the rapid growth of the fund, reflecting partly the income from natural resources in Figure 1, but also the high rates of return the fund has enjoyed since inception. By the end of 2020, the fund was around 3.2 times GDP, with a clear upward trend. Figure 3 shows the average share of wealth in risky assets over time, which is mostly in internationally diversified equity investments. The average allocation in stocks was around 40 percent in the first decade and has gradually risen to around 70 percent (with a small allocation in real estate) by 2018, and 72.4 percent in 2021.

There are at least three main motivations for starting a SWF for any country. The first is intergenerational: the ability of a country to save both for current and future generations. The second is to smooth shorter-run variations in the business cycle. Both of these motives are explicitly stated in the strategic objectives of Norway's SWF and that is the reason we incorporate them explicitly in the normative objective function we use.

The third motivation is relatively more controversial and relates to the need to avoid problems of competitiveness arising from a rising real effective exchange rate if large natural resources proceeds are invested domestically (commonly known as the "Dutch disease" after Holland's experience in the 1950s). Figure 4 plots the ten year rolling average real GDP growth in Norway since 1980. The graph points to a declining trend in the growth rate, even after the establishment of the SWF in 1996. We plot this annual GDP growth in Figure 4 because the trend GDP growth in an economy can prove to be an important determinant of SWF accumulation in the context of smoothing risks across generations. The rapid drop from around 4.5 percent in GDP growth in the early 1980s to close to 1 percent in 2020 means that this variable can prove important in making decisions with regards to optimal SWF management and our model can provide insights into how such changes should be affecting optimal decisions in the fund with regards to accumulation, withdrawals and asset allocation.

3 The Model

3.1 Preliminary assumptions

A key issue in the analysis involves the specification of the relevant objective function. To proceed with the normative analysis we make the assumption that the sovereign has dealt with any corporate governance or political issues that might arise in SWF management. The literature emphasizes that substantial heterogeneity in potential additional objectives exists (like gaining political influence locally and internationally) and corporate governance quality might also vary across SWFs (see Bernstein et al. (2013)). There is also evidence for fiscal procyclicality in many resource-dependent economies (see Coutinho et al. (2022) and references therein). We view the pure economic problem facing a sovereign sufficiently interesting and complex and therefore focus only on that part of normative SWF management (designing strategic asset allocation rules for a "benevolent SWF"). The government is assumed to care about a proxy for government consumption expenditures (C_{gt}) that need to cover important commitments that are decided upon through a political process. Before explicitly writing down the objective function of the government, we describe the economic environment and the available choices to the government.

3.2 Business Cycle

We assume that Y_t, the state of the domestic business cycle, is well captured by a deterministically detrended series of GDP that is persistent. This allows us to capture the experience of many countries around the world that have varying growth experiences. We therefore assume that detrended Y_t follows an AR(1) process

$$Ln(Y_{t+1}/\overline{Y}_{t+1}) = \rho \cdot Ln(Y_t/\overline{Y}_t) + \varepsilon^Y_{t+1},\ \varepsilon^Y_{t+1} \sim N(0, \sigma^2_{\varepsilon_Y}). \tag{1}$$

$$\overline{Y}_t = \overline{Y} \cdot e^{gt} \tag{2}$$

Deterministic growth is captured by g. The experience of Norway from Figure 4 illustrates already that this might be a key variable to consider given the preference for intergenerational equity. The volatility and persistence of detrended GDP and the correlation of its innovation ε_Y with other exogenous background risks will also be important parameters. They will be the subject of an extensive calibration and be central in devising the normative implications for the SWF. Indeed, many developing countries design SWFs because they face high GDP growth volatility and also trend growth (see Aguiar and Gopinath (2007) for evidence on emerging market business cycles). We also normalize $\overline{Y}$ to be equal to one.

3.3 Natural Resources

Each year the value, in real terms, of the natural resources that are extracted is given by

$$P_t Q_t \tag{3}$$

where Q_t is the real quantity extracted and P_t is the (inflation-adjusted) price of the natural resource.

In the baseline model we assume:

- A linear (exogenous) extraction rate from zero in the first year to a certain value.

- A deterministic date T^* at which the production of the commodity begins to decrease and falls to zero at different speeds.

- A function $f(t - T^*)$ that models the exogenous depletion of the resource. Mathematically, Q_t becomes

$$Q_t = \begin{cases} Q_t & \text{for } T \leq T^* \\ f(t - T^*) & \text{for } T > T^* \end{cases} \tag{4}$$

We experiment with different functions but the baseline will be one where the resource completely disappears after T^*, that is, $f(t - T^*) = 0$ for $T > T^*$. We abstract from endogeneizing the extraction process and leave that extension to future work.

Moreover, we assume that P_t follows an AR(1) process in logs (use lower case for logs)

$$p_{t+1} = \kappa_p \cdot (1 - \rho_p) + \rho_p \cdot p_t + \varepsilon^p_{t+1},\ \varepsilon^p_{t+1} \sim N(0, \sigma^2_{\varepsilon^p}). \tag{5}$$

This assumption is consistent with the empirical evidence in Deaton and Laroque (1992, 1996) who show that, at an annual frequency over the twentieth century, many real commodity prices are well captured by an AR(1) model. We use our own analysis to calibrate this process on an international commodity index later on, further confirming this assumption.

3.4 Financial Assets

The investment opportunity set is assumed to be constant and there are up to three financial assets:

- Riskless asset (treasury bills or cash) with real gross return, $R^f = e^{r^f}$

- Stocks with real gross return given by (use lower case for log returns)

$$r^S_{t+1} - r^f = \mu + \varepsilon^S_{t+1},\ \varepsilon^S_t \sim N(0, \sigma^2_{\varepsilon^S}) \tag{6}$$

We allow for correlation between (real) stock returns and (real) price of the natural resource ($corr(\varepsilon^S_t, \varepsilon^p_t)$).

The SWF portfolio return is then given by

$$R^p_{t+1} = \alpha^S \cdot (R^S_{t+1} - R_f) + R_f \tag{7}$$

3.5 Budget Constraint

Each period the SWF's financial wealth evolves as follows,

$$W_{t+1} = R^p_{t+1} \cdot (1 - k_t) \cdot W_t + P_{t+1} \cdot Q_{t+1} \quad (8)$$

$$0 \leq k_t \leq 1 \quad (9)$$

where $R^p_{t+1} = 1 + r^p_{t+1}$ is the SWF portfolio return and $P_t Q_t$ is the value of the natural resource extracted in year t. The part of the wealth that is being extracted every period is denoted by k_t. Unlike the exogenous evolution of the value of the commodity, k_t will be a key endogenous choice for the SWF. The different implications emanating from different choices of k_t will be discussed below.

3.6 Objective Function

We assume the government's objective function is given by

$$V = MaxE_0 \sum_{t=0}^{T} \beta^t \cdot U(C_{gt}), \quad (10)$$

where $C_{gt} = \theta \cdot Y_t + D_t - h \cdot \overline{Y}_t$, $U(C_{gt}) = \frac{C_{gt}^{1-\gamma}}{1-\gamma}$, γ is the coefficient of relative risk aversion, β is the discount factor, D_t is the payout from the SWF (a dividend or withdrawal from the fund), and h is the minimum subsistence fraction of trend output that needs to be covered by the payouts (government commitments).

We can express D_t as a fraction of the wealth in the SWF, so

$$D_t = k_t \cdot W_t. \tag{11}$$

We do that because we will assume the government is trying to withdraw a fraction of accumulated financial wealth from the SWF so that on average the fund does not get depleted. To do that, we will assume the fraction k_t is (the volatility-adjusted and growth-adjusted) expected annuity from accumulated wealth, that is, the government can withdraw the net expected return each period. Specifically,

$$k_t = E_t r^p_{t+1} \tag{12}$$

It is important to discuss these choices for the SWF's objective function. First, there is a preference for smooth payouts from the SWF to reduce background uncertainty, given the concave utility function. The payout (D_t) is an endogenous choice and this preference is consistent with the strategic goal of smoothing the benefit across different generations; we make the additional assumption that there is a preference not to decumulate wealth through the expected annuity. In the words of Campbell and Sigalov (2022) this is an "arithmetic sustainable spending constraint model" that will require consumption and wealth to remain constant over time after the natural resource disappears, but will also imply a constraint on how much the SWF can be relied upon to smooth business cycle fluctuations while the resource lasts.[5] This assumption can be relaxed later to be compared to the case where the

[5]The results for a "geometric sustainable spending constraint model" using the adjustments proposed by Dybvig and Qin (2021) will be qualitatively similar to the results under the "arithmetic sustainable spending

withdrawal is not the expected annuity value but instead chosen optimally, or chosen under specific constraints.

Second, we introduce an exogenously evolving process given by $\theta \cdot (Y_t)$, where Y_t is the exogenously evolving GDP in the economy. The parameter θ captures the amount of government consumption that can be financed by tax revenues. Tax revenues directly rise from higher GDP. As a result, depending on the value of θ, tax revenues will be higher (lower) during expansions (recessions). This exposure will naturally affect the choice of how much to withdraw and the optimal asset allocation.

The natural question is how to calibrate a reasonable range for the parameter θ. If θ=0, then the sovereign cannot rely on tax revenues to finance government consumption at all. This is an extreme case but might still instructive to investigate what the model implies in that instance. Countries that are really poor and suddenly discover natural resources might be in this category, for example. When θ is positive, the state of the economy can affect both the payout and optimal asset allocation in different ways. In a boom (recession), when the deviation from trend GDP is positive (negative), there might be less (more) urgency to give a payout, while asset allocation might be affected by the ability to take on a more (less) risky position in equities. The business cycle state variable will be modelled as persistent, to a first approximation capturing the growth rate of the economy, and therefore will be an exogenous state variable.

Another component of the objective functions is the minimum fraction of trend growth

constraint model" because they involve reducing the withdrawal rate by a certain constant fraction of wealth. This will imply a higher wealth accumulated by the SWF but will not change the qualitative predictions of the model. The quantitative predictions will change depending on the size of the adjustment but the comparative statics we report will continue to hold.

GDP (h) that acts as a habit, since it is a minimum commitment that the SWF needs to satisfy period by period relative to trend growth. This can also be set to zero but higher values can be thought of as higher risk aversion when the SWF approaches the pre-committed expenditures a government or a university has already made to public sector workers or faculty, respectively. Future work can experiment with a time-variation in this parameter.

Finally, the horizon is important. We assume a finite horizon so as to allow us to potentially consider the case of $\beta = 1$, but we will take T, the horizon of the sovereign, to be very large (up to 500 years) so that it is representative of a very long-lived problem. Both the horizon length and the high discount factor capture the idea that SWFs are intended for future generations.

Given the deterministic growth in the economy, we need to detrend the model to make it stationary. Define lower case variables that require detrending by GDP as upper case variables divided by the deterministic trend in GDP (that is, $q_t = Q_t/\overline{Y}_t$).[6] We then have the normalized value function as follows

$$v_t(w_t, p_t, y_t) = \underset{\{\alpha_t\}}{MAX}\, U(c_{gt}) + \beta \cdot e^{g(1-\gamma)} \cdot E_t v_{t+1}(w_{t+1}, p_{t+1}, y_{t+1}) \tag{13}$$

subject to

$$c_{gt} = \theta \cdot y_t + k_t \cdot w_t - h \tag{14}$$

$$w_{t+1} = e^{-g} \cdot R^p_{t+1} \cdot w_t - e^{-g} \cdot R^p_{t+1} \cdot k_t \cdot w_t + e^{p_{t+1}} \cdot q_{t+1} \tag{15}$$

$$k_t = E_t r^p_{t+1} \tag{16}$$

[6]Returns and the real commodity price are not detrended by GDP.

$$ln(y_{t+1}) = \rho \cdot ln(y_t) + \varepsilon^Y_{t+1} \tag{17}$$

and the exogenous stochastic processes for commodity prices and stock returns that require no normalization.

We investigate two main different versions of this problem to understand better the implications for SWF management emanating from different assumptions. We summarize them as follows:

1) The withdrawal is a function of asset allocation when the commodity revenues are present, and reverts to an annuity after the resource is exhausted.

2) The withdrawal is chosen optimally when the commodity revenues are present, and reverts to an annuity after the resource is exhausted.

In both variants, we can add additional constraints like the exogenously imposed fiscal rule used by the Norwegian SWF (that withdrawals cannot be greater than a certain fraction (say 4 percent) of the accumulated wealth) or that the asset allocation to stocks has to be between certain bounds.

4 Calibration

4.1 Preference parameters

We view the SWF as a very patient entity with a long horizon. We therefore use a discount factor for an annual frequency model equal to 0.98. We also view the SWF as risk averse, reflecting the risk aversion of government authorities or politicians having career concerns acting in a risk averse way: we therefore use a value of 4.0 for risk aversion.[7]

A key parameter is θ. For a country like the U.S. the average share of government expenditures to GDP from 1929 to 2018 is 20 percent and tax revenues are around 17 percent. We therefore use 0.2 as a baseline and experiment with 0.1 and 0.3. This parameter can vary depending on the state of tax capacity for the country starting the SWF; we provide sensitivity analysis with respect to this parameter.

4.2 Commodity Price Process

Based on our reading of the empirical evidence and our own calculations based on an world commodity index provided by the IMF, we estimate an AR(1) with a coefficient (ρ_p) equal to 0.7 and an unconditional standard deviation equal to 15 percent to calibrate the background risk in the commodity prices. This is consistent with the annual estimates provided by Deaton and Laroque (1996) as a firxt approximation to the highly non-linear processes followed by real commodity prices at this frequency.

4.3 Extraction

We assume an exogenous extraction process that can be augmented depending on the experience of each country. For the purposes of this paper, we assume extraction starts at 5 percent of GDP and linearly increases to 15 percent after 25 years. It stays constant at that point for another 50 years and then is linearly depleted from year 75 to year 85 where it reaches zero. We experiment with different exogenous extraction processes but we do not

[7] Comparative statics results from changing the discount factor and risk aversion preference parameters are available upon request.

report these experiments in the paper as this calibration can vary across countries, and the model can be readily adjusted to reflect different assumptions (we have experimented with non-linear profiles as well where the extraction happens at a higher pace in the beginning).

4.4 Horizon

The SWF will have a mandate that is presumably given to it by the relevant authority (the ministry of finance, for example). To maintain discipline in the experiments, we keep the extraction horizon constant but we make different assumptions on the horizon of the SWF. Our baseline horizon (T) in Equation (10) is 500 years, which means the mandated horizon is around five times longer than the expected depletion of the natural resource.

4.5 Return processes

We assume that the SWF invests in an internationally well diversified stock market index and a risk free asset. The Vector Autoregression (VAR) we use during the extraction of the commodity is made up of the real commodity price, the detrended GDP and the stock return. We experiment with different parameters of the variance-covariance matrix associated with this VAR.

The annual risk free rate is 1 percent, while the equity risk premium is 5 percent with an annual standard deviation equal to 20 percent. The mean growth in the economy is zero percent, while the standard deviation of detrended GDP is 5 percent, given the higher volatility of countries with large natural resources, and the tendency for the cycle to be more important than the trends in many of these countries (for example, Aguiar and Gopinath

(2007)). The autocorrelation coefficient for detrended GDP is set at 0.3.

In terms of correlations, we start with a zero correlation between the innovations in stock returns and the commodity price, the correlation between stock returns and detrended output is also zero and the correlation between detrended output and the commodity price is equal to 0.25. We provide comparative statics results later to quantify the hedging demands arising from these correlations.

4.6 Solution Method

We solve the problem backwards using a relatively standard stochastic dynamic programming approach. We need to make an assumption on what the SWF should be doing after the natural resource is depleted. A natural assumption is to consume the expected return, so that the SWF wealth is never depleted. We also impose a smooth-pasting condition by guessing and verifying the share of wealth in stocks during this period that would provide a smooth transition from the positive resources period to the zero resources one.

5 Normative results

5.1 Baseline model: A simulation

We start by picking one typical simulation from the baseline model at random and plotting the main exogenous shocks and endogenous outcomes in Figure 5. Panel A shows the randomness around the typical mean of the value of the extracted commodity as a percent of GDP. The uncertainty reflects the realized values of the exogenously simulated commodity price in Panel B. We have assumed away uncertainty in extraction, therefore around the mean profile of the extracted commodity, the value fluctuates based on the internationally traded (and therefore exogenous) commodity price. In periods between around 30 and 40, when the commodity price is low, the value of the extracted commodity has a pronounced and persistent drop. By assumption, after 85 years, the extracted commodity drops to zero. The third exogenous shock is the detrended GDP which is shown in Panel C.

The other three panels show the normative implications of the model. The simulated wealth to GDP (Panel E) reaches its maximum value when the natural resource is completely depleted after 85 years. Simulated withdrawals in GDP units are relatively substantial given the profile of the extracted value in Panel A. Simulated withdrawals (Panel F) track simulated wealth given the sustainability constraint. Given that the extracted value of the natural resource is quite high, there is substantial wealth that accumulates in the SWF; this can reach 20 GDP units in just 85 years, which implies a substantial withdrawal level as well (60 percent of GDP at the highest point). The SWF takes aggressive asset allocation positions in the early years (Panel D) and gradually reduces the exposure as the time of resource depletion approaches (in this case after around 50 years).

5.2 Long Run: Average simulations

The previous subsection picked a random simulation to illustrate what might happen in the context of the model. Averaging across all possibilities can give us a better understanding of the SWF's strategic asset allocation. Figure 6 does so illustrating over six Panels the average behaviour of the three key endogenous decisions (withdrawals (Panel A), wealth (Panel B) and share of wealth in stocks (Panel C)). Figure 6 also shows what happens when we move from a zero to a one percent deterministic GDP growth rate.

Panel A includes the mean value of the extracted commodity as a percent of GDP. Given the imposed sustainability constraint, the average withdrawal inherits the shape of mean wealth. Panel C shows the preference for taking more risk early on in the life of the SWF, and derisking the portfolio as resource depletion date approaches. This result arises from the same intuition that most life-cycle models (for example, Cocco et al. (2005)) predict higher stock exposure early on in the life cycle. The background risk in this case is the commodity price which is not correlated with stock market risk, in the same way idiosyncratic labor income risk is typically uncorrelated with aggregate stock market risk in the life-cycle literature. The commodity is like an implicit riskless asset as a result, and therefore the financial portfolio is invested in equities early on. As the implicit riskless asset in the form of the commodity is depleted, riskless assets are added to the portfolio.

Figure 6A compares the results of the baseline model where aggregate growth is zero to a case where aggregate growth is one percent per year. The higher growth allows the SWF to take on higher stock market risk (Panel C), and through the sustainability extraction constraint simulated average withdrawals are equal or higher than the respective case in the baseline model (Panel A). Strikingly, the simulated average wealth is the same in both cases, as the higher average allocation in stocks generates a higher average expected return that results in higher average withdrawals, leaving average simulated wealth unchanged (Panel B).

Figure 6B provides the range of possibilities that might arise by plotting the same variables between their 10th and 90th percentiles for the baseline model. It can be seen that a relatively wide range of outcomes is possible given the large uncertainty faced by the SWF. The range of withdrawals might be between zero and 60 percent of GDP depending on the time since SWF inception and the background uncertainty. The same applies for the size of the SWF; it can range at year 85 (when the resource is depleted) between around 4 GDP units to around 20 GDP units for the 10th and 90th percentile, respectively.

5.3 Short Run: Effect of Commodity Prices

Figures 7A and 7B compare policy functions to analyze the results when we condition on the mean detrended GDP for the baseline correlation between the GDP and commodity price innovations (0.25 in Figure 7A) and a correlation that is much higher (0.9 in Figure 7B). The much higher correlation is justified for a country exporting a natural resource that makes up a large component of GDP. A country like Saudi Arabia would be an appropriate example; when world oil prices are high, oil revenues are high and the correlation between GDP and oil prices will also be very high.

Figure 7A shows withdrawals (left column) and the share of wealth in stocks (right column) as a function of the wealth state variable. Each row is for different times since SWF inception (25, 55 and 100 years). The striking feature is that the commodity price realization does not quantitatively affect portfolio allocation at any horizon, and this is then reflected in the withdrawal policy.

The picture changes in quantitatively important ways when the correlation between the commodity price and GDP innovations increases to 0.9 (Figure 7B). In this instance, it can be seen that as long as the natural resource is getting extracted (years 25 and 55 in the first two rows), a low commodity price realization implies a higher asset allocation in stocks (Panel D or E), which leads to a higher withdrawal rate (Panel A or B), for a given wealth. Low commodity prices are associated with below trend GDP levels and therefore stabilization motives imply higher withdrawals and higher allocations to the stock market. The effect disappears when the natural resource is depleted (Panels C and F), as expected. This prediction of taking more risk in bad times arises jointly from the sustainability constraint and the stabilization motive in the objective function.

5.4 Short Run: Effect of the Business Cycle

We now repeat the same experiments but investigate the effect of the business cycle when the commodity price is at its mean level. Figure 8A and 8B present the same policy functions as in Figures 7A and 7B but we now condition on mean commodity prices and vary the realizations of the GDP deviation from trend. In periods where GDP is below trend and the natural resource still exists (25 and 55 years after inception), the withdrawal is higher and the asset allocation in stocks is higher, for a given level of cash on hand. The stabilization need is higher in recessionary states and therefore a higher withdrawal is needed, which implies a higher allocation to the stock market. The effect disappears when the natural resource gets depleted (year 100 in Panels C and F).

Figure 8B repeats the same experiment but now the correlation between the GDP and commodity price innovations is much higher (0.9). The direction of the effects qualitatively is identical to Figure 8A but the quantitative magnitudes are much weaker. A stronger correlation between commodity prices and GDP implies a much lower reaction of the asset allocation and withdrawal rate as a response to the state of the business cycle. It is basically harder to hedge GDP shocks and the withdrawal and associated porftolio responses are much weaker as a result.

5.5 Varying Taxation Capacity

In choosing the SWF's objective function, Equation (14) adds to what the government is trying to smooth a parameter (θ) multiplying (detrended) GDP. There are different interpretations of this parameter. Our preferred interpretation is that the SWF should be less relied upon when GDP is high, because tax revenues are naturally higher in those states and/or government expenditure funding needs lower, and this is a natural way to introduce this idea. We can investigate special cases around this parameter, however, since that actual choice might differ across countries based on their state of development and their political preferences. We therefore investigate special cases when $\theta = 0$ and the local business cycle therefore does not interfere with SWF management.

Figure 9 shows average simulated results for the cases, $\theta = 0.0, 0.2$. The main lesson is that asset allocation becomes a lot more conservative when $\theta = 0$ (see Panel C), and therefore at the same time there are fewer withdrawals on average (Panel A) throughout the life of the SWF. The ability to rely on tax revenues therefore allows the SWF to take more risky positions. Essentially, in the absence of a minimum (exogenous) buffer provided by local GDP, the SWF should be run more conservatively. Strikingly the level of average SWF accumulation stays the same in either case; the more conservative portfolio generates lower withdrawals and average wealth stays unchanged (Panel B).

5.6 Low real interest rates

We have witnessed low interest rates in the last decade but also a trend towards lower real interest rates globally since the 1990s. Campbell and Sigalov (2022) point out that static models with sustainable spending constraints can generate asset allocation behavior that resembles "reaching for yield" investment behavior with higher exposure in the stock market as interest rates decline to maintain a certain level of (sustainable) income (withdrawals).

Figure 10 investigates this effect in the context of our SWF model when we vary the risk free rate from 1 percent to zero to negative one percent, ceteris paribus. Panel C shows that there are quantitatively important shifts in the asset allocation when the SWF approaches maturity. In our case and for this calibration this occurs after around 30 years, and the effect gets magnified and remains quantitatively large for ever. By year 85 when the natural resource is depleted, the asset allocation moves from around 40 percent to around 70 percent in stocks when the interest rate moves from one to negative one percent. Despite the much higher asset allocation in stocks in the lower interest rate environment, the average withdrawal rate remains lower (Panel A). Interestingly, the combination of Panels A and C means that the average accumulated wealth stays approximately the same across the different interest rate environments.

5.7 Different Constraints on Withdrawals

Many SWFs impose different constraints on withdrawals. In the Norwegian SWF, the maximum withdrawal rate is typically 4 percent, for example. We investigate the implications of two constraints in this subsection. We call the first a k-capped constraint where the withdrawal rate has to be less than or equal to a four percent return from the SWF. This constraint is intended to mirror the Norwegian constraint. A second possibility is an α-capped constraint where the maximum allowable share of wealth in stocks is sixty percent. This is intended to capture the idea that many governments with young SWFs might be uncomfortable investing all the proceeds in international capital markets and/or might impose constraints to invest large parts of the proceeds domestically.[8]

Figure 11 compares the implications of the different constraints. Starting from Panel D, we note that the k-capped constraint generates a similar shape in portfolio choices as the baseline model. Nevertheless, the average withdrawal rate is constrained at 4 percent (Panel B), whereas in the baseline model the withdrawal rate is unconstrained and is higher (six percent). This would be hidden if we did not normalize withdrawals by wealth (Panel A). As a result, even though asset allocations are essentially identical between the baseline model and the k-capped model, the withdrawal rate relative to accumulated wealth is lower, and therefore there is a higher accumulated wealth in the k-capped model (Panel C). This leads eventually to a higher withdrawal rate (Panel A). We conclude that the comparative statics we have performed will have similar shape in the k-capped economy as in the baseline economy.

[8] There is evidence for such behavior for pension funds from the U.S. (Bradley et al. (2016)) and it should be expected that many governments might have similar political biases.

Imposing a constraint on the average share of wealth in stocks affects primarily the asset allocation decision. As can be seen in Panel D, the shape is similar as in the baseline economy but the upper bound is kept at sixty percent rather than reaching 100 percent in the early life of the fund. The same derisking policy is followed as the natural resource is depleted, but this begins to happen much later. As a result of the lower share of wealth in stocks, average withdrawal rates are also lower than in the baseline (Panel A) and this implies that the average wealth accumulation remains the same as in the baseline case (Panel C).

5.8 Habit (Expenditure Commitments)

We next investigate the role of the additive habit (intended to portray expenditure commitments) in our model. We do not model a habit as a state variable that could be evolving either endogenously or exogenously to avoid having an additional state variable in the model. We instead follow the simpler approach of investigating what happens if the SWF also cares about smoothing not only the level of the government consumption but also its level relative to the deterministic GDP trend.

We can see from Figure 12 that the SWF acts in a more risk averse way in the presence of this habit (Panel C). The average share of wealth in stocks is below the baseline in all time periods. As a result, the average withdrawal rate is less than in the baseline (Panel A). Interestingly, this combination again makes the average wealth being accumulated the same across both specifications. Our conclusion is that this commitments undertaken by a SWF (the equivalent of the habit in the model) can have a substantial effect on asset allocation and withdrawal rates and could be a good topic for future research in terms of further understanding how the commitments evolve and when they are more important and binding as that can affect ex ante asset allocations.

6 Alternative Formulation: Optimal Withdrawals

The choice of the objective function is key and one can think that different political processes and opinions can result in different given mandates to an independent SWF. One key question is whether, in the face of expected secular growth, the SWF should be withdrawing from the fund using the constraint on expected returns. What happens if the fund is instead allowed to optimally choose the withdrawal rate without any constraints, while the natural resource exists?

Mathematically, the problem is identical to the one solved up to now with the exception that the withdrawal rate is chosen optimally. Figure 13 shows the results from this extension. The main difference from the baseline is that when the SWF is young, the withdrawal rate is much higher; the optimal withdrawal rate tracks the resource revenues in the first ten years. This is to be expected from an entity that is expecting higher resource revenues in the future and is not fully patient (β is not equal to one). As a result of the higher withdrawal rate, mean wealth is lower than in the baseline model and this generates a higher average allocation to stocks due to the usual intuition of natural resource revenues being viewed as an implicit riskless asset; when less financial wealth is accumulated but the implicit riskless asset value is the same, more risky investments are made. Because of the higher asset allocation in stocks for a large part of the accumulation phase, but also after the resource ends, average withdrawals are higher than in the baseline model after period 85 (end of the resource).

What happens when there is a positive expected growth rate in the economy? The comparisons are shown in Figure 14. The higher expected growth rate acts as a higher discount rate when optimal withdrawals are allowed. A permanently richer government can afford to withdraw more aggressively as the fund matures and only starts accumulating wealth after around 20 years and at a very slow pace. The total accumulated wealth is much lower than in the baseline sustainable-spending model (Panel B), and as a result a higher share of wealth is invested in stocks (Panel C). This is essentially the same intuition as in life-cycle models of consumption (for example, Carroll (1997)).

This analysis emphasizes the importance of understanding the implications of different SWF objective functions. An economy with a higher expected growth rate, for example, that does not take into account a sustainability constraint in the withdrawal rates should have a much smaller SWF. A political process that needs to arrive at a mandate for an independently-managed SWF could reach different decisions once the implications of different objective functions are better understood.

7 Conclusion

We have proposed and solved a normative model of strategic asset allocation for a SWF that aims to balance the trade-off between short run stabilization and long run intergenerational equity in the presence of a sustainable-spending constraint. The background risk arising from commodity prices implies a more aggressive investment strategy in the early years of the fund with portfolio derisking as the resource extraction approaches the end. Choosing withdrawals optimally as opposed to through a sustainable-spending constraint substantially changes the predictions of the model with regards to the SWF size, and optimal asset allocation, especially in the presence of higher expected economic growth. The government's ability to rely on taxation capacity to smooth aggregate shocks or the presence of consumption commitments imposed on the SWF also substantially alter asset allocation. Moreover, a low real interest rate environment generates a natural "reaching for yield" aggressive asset allocation strategy in the presence of the sustainable-spending constraint.

References

Aguiar, M. and G. Gopinath (2007). Emerging market business cycles: The cycle is the trend. *Journal of Political Economy 115*(1), 69–102.

Ang, A. (2014). *Asset management: A systematic approach to factor investing.* Oxford University Press.

Ang, A., M. Brandt, and D. Denison (2014). Review of the active management of the norwegian government pension fund global. *Report to the Norwegian ministry of finance*.

Bernstein, S., J. Lerner, and A. Schoar (2013). The investment strategies of sovereign wealth funds. *Journal of Economic Perspectives 27*(2), 219–38.

Bradley, D., C. Pantzalis, and X. Yuan (2016). The influence of political bias in state pension funds. *Journal of Financial Economics 119*(1), 69–91.

Campbell, J. Y. and R. Sigalov (2022). Portfolio choice with sustainable spending: A model of reaching for yield. *Journal of Financial Economics 143*(1), 188–206.

Campbell, J. Y. and L. M. Viceira (2002). *Strategic Asset Allocation.* Oxford University Press.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics 112*(1), 1–55.

Chetty, R. and A. Szeidl (2007). Consumption commitments and risk preferences. *The Quarterly Journal of Economics 122*(2), 831–877.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Cochrane, J. H. (2022). Portfolios for long-term investors. *Review of Finance 26*(1), 1–42.

Coutinho, L., D. Georgiou, M. Heracleous, A. Michaelides, and S. Tsani (2022). Limiting fiscal procyclicality: Evidence from resource-dependent countries. *Economic Modelling 106*, 105700.

Deaton, A. and G. Laroque (1992). On the behaviour of commodity prices. *The Review of Economic Studies 59*(1), 1–23.

Deaton, A. and G. Laroque (1996). Competitive storage and commodity price dynamics. *Journal of Political Economy 104*(5), 896–923.

Dybvig, P. H. and Z. Qin (2021). How to squander your endowment: Pitfalls and remedies. *Available at SSRN 3939984*.

Gilbert, T. and C. Hrdlicka (2015). Why are university endowments large and risky? *The Review of Financial Studies 28*(9), 2643–2686.

Kotter, J. and U. Lel (2011). Friends or foes: Target selection decisions of sovereign wealth funds and their consequences. *Journal of Financial Economics 101*(2), 360–381.

Viceira, L. M. (2001). Optimal portfolio choice for long-horizon investors with nontradable labor income. *Journal of Finance 56*, 433–470.

Figure 1: Norway's natural resource income as a proportion of GDP (1970-2019)

Source: The Global Economy.com, The World Bank

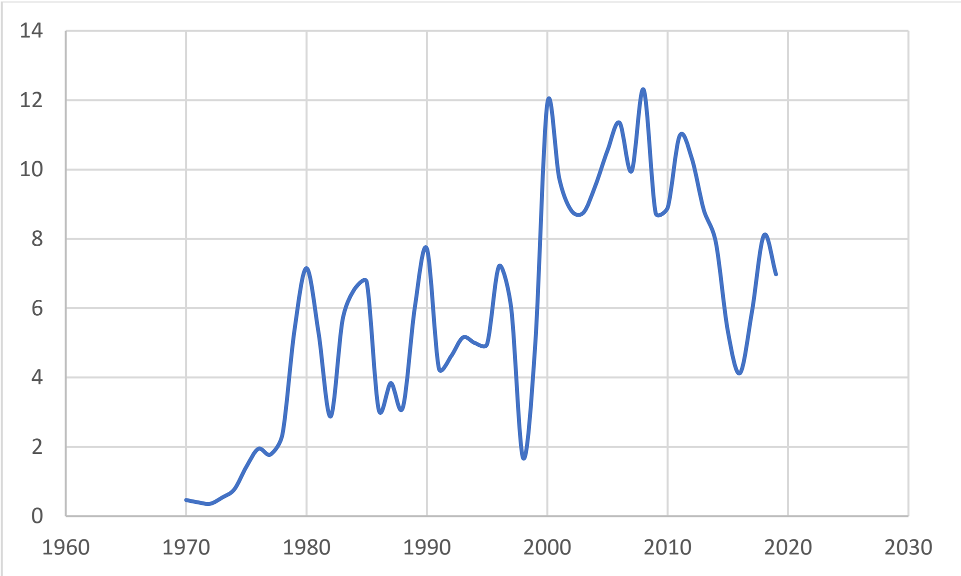

Figure 2: Size of Norway's Pension Fund Global relative to GDP

Source: https://www.nbim.no/en/the-fund/market-value/ and Statistics Norway. SWF established in 1996.

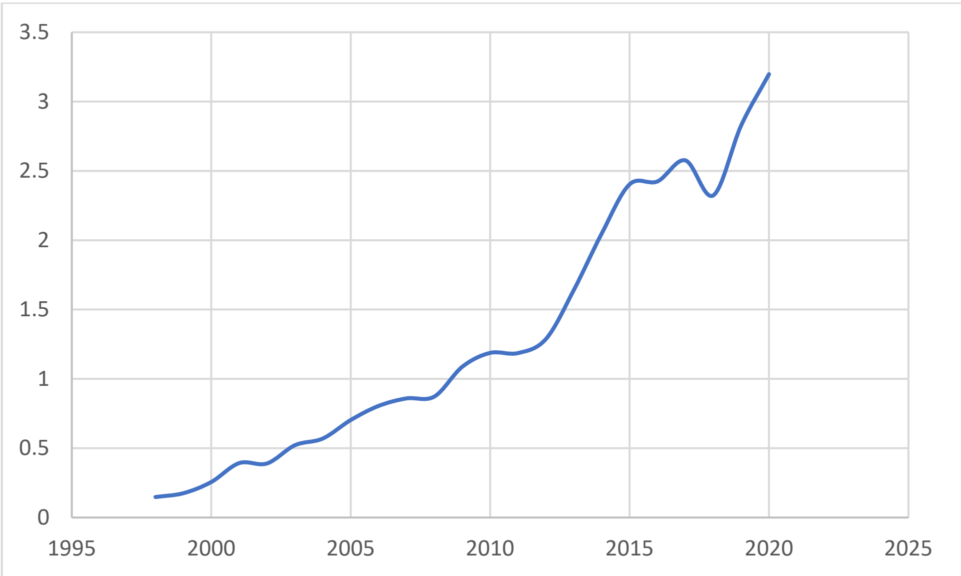

Figure 3: Asset Allocation for Norway's Pension Fund Global

The graph shows the risky assets as a share of the total SWF value. Risky assets refer to stocks and real estate over total wealth and the subset is stocks over total wealth. Real estate became positive in 2011 but remains less than five percent of the total asset allocation. Source: https://www.nbim.no/en/the-fund/market-value/.

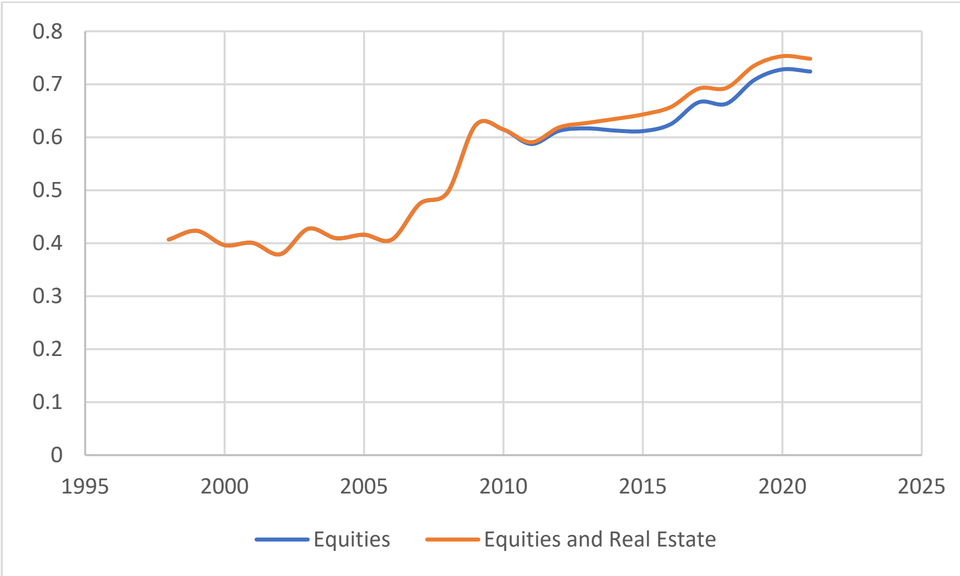

Figure 4: Real Per Capita GDP Growth in Norway

The graph shows the 10-year rolling average real GDP growth rate in Norway. Source: Statistics Norway.

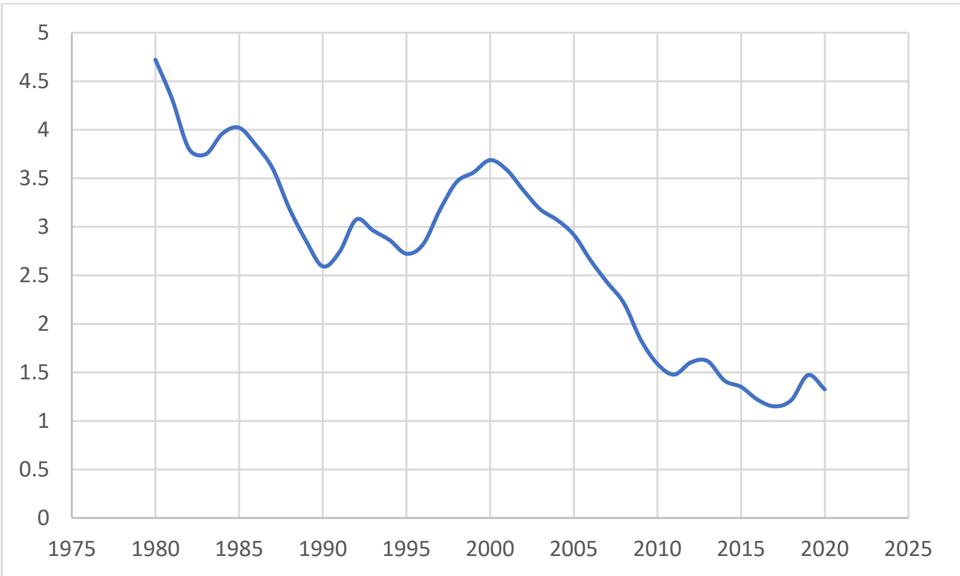

Figure 5: Typical Simulation for Baseline Model

Figure 5 presents results for a typical simulation from the baseline model where β=0.99 and μ_g is 0%. Panels A, B and C display the exogenous uncertainty and Panels D, E and F the endogenous responses.

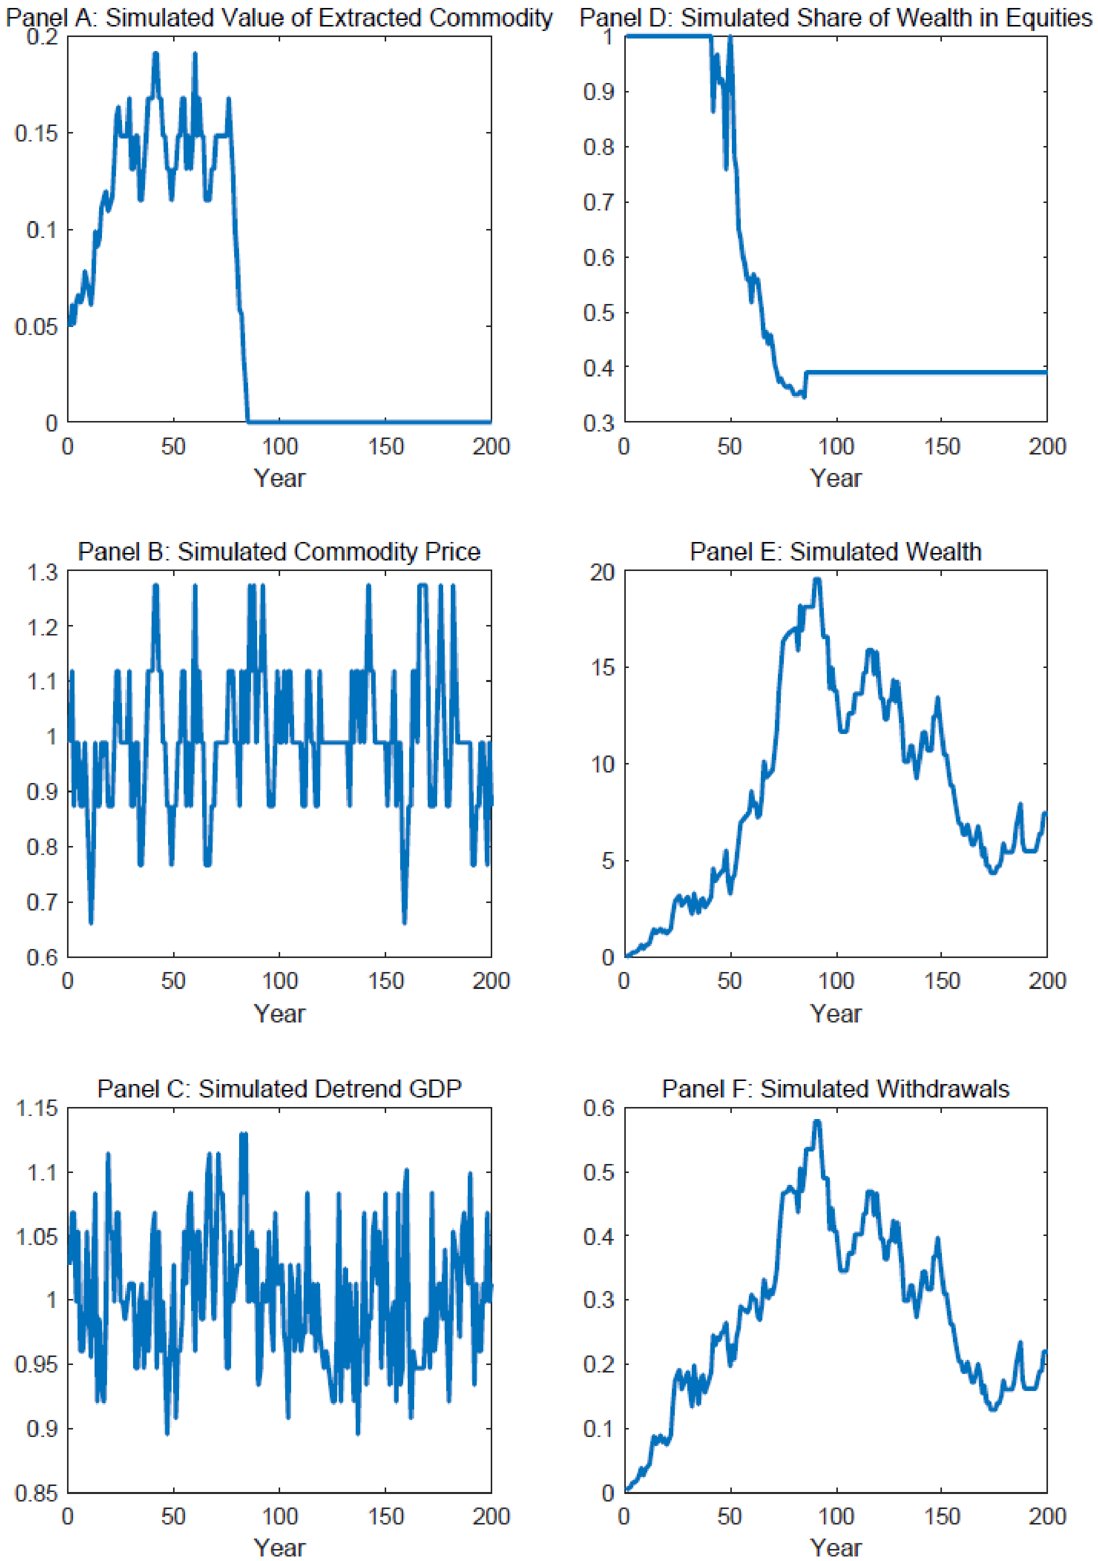

Figure 6A: Average Simulations for Baseline Model

Figure 6 presents results from the baseline model where β=0.98, no habits, r_f=1% and μ_g is 0 and 1%.

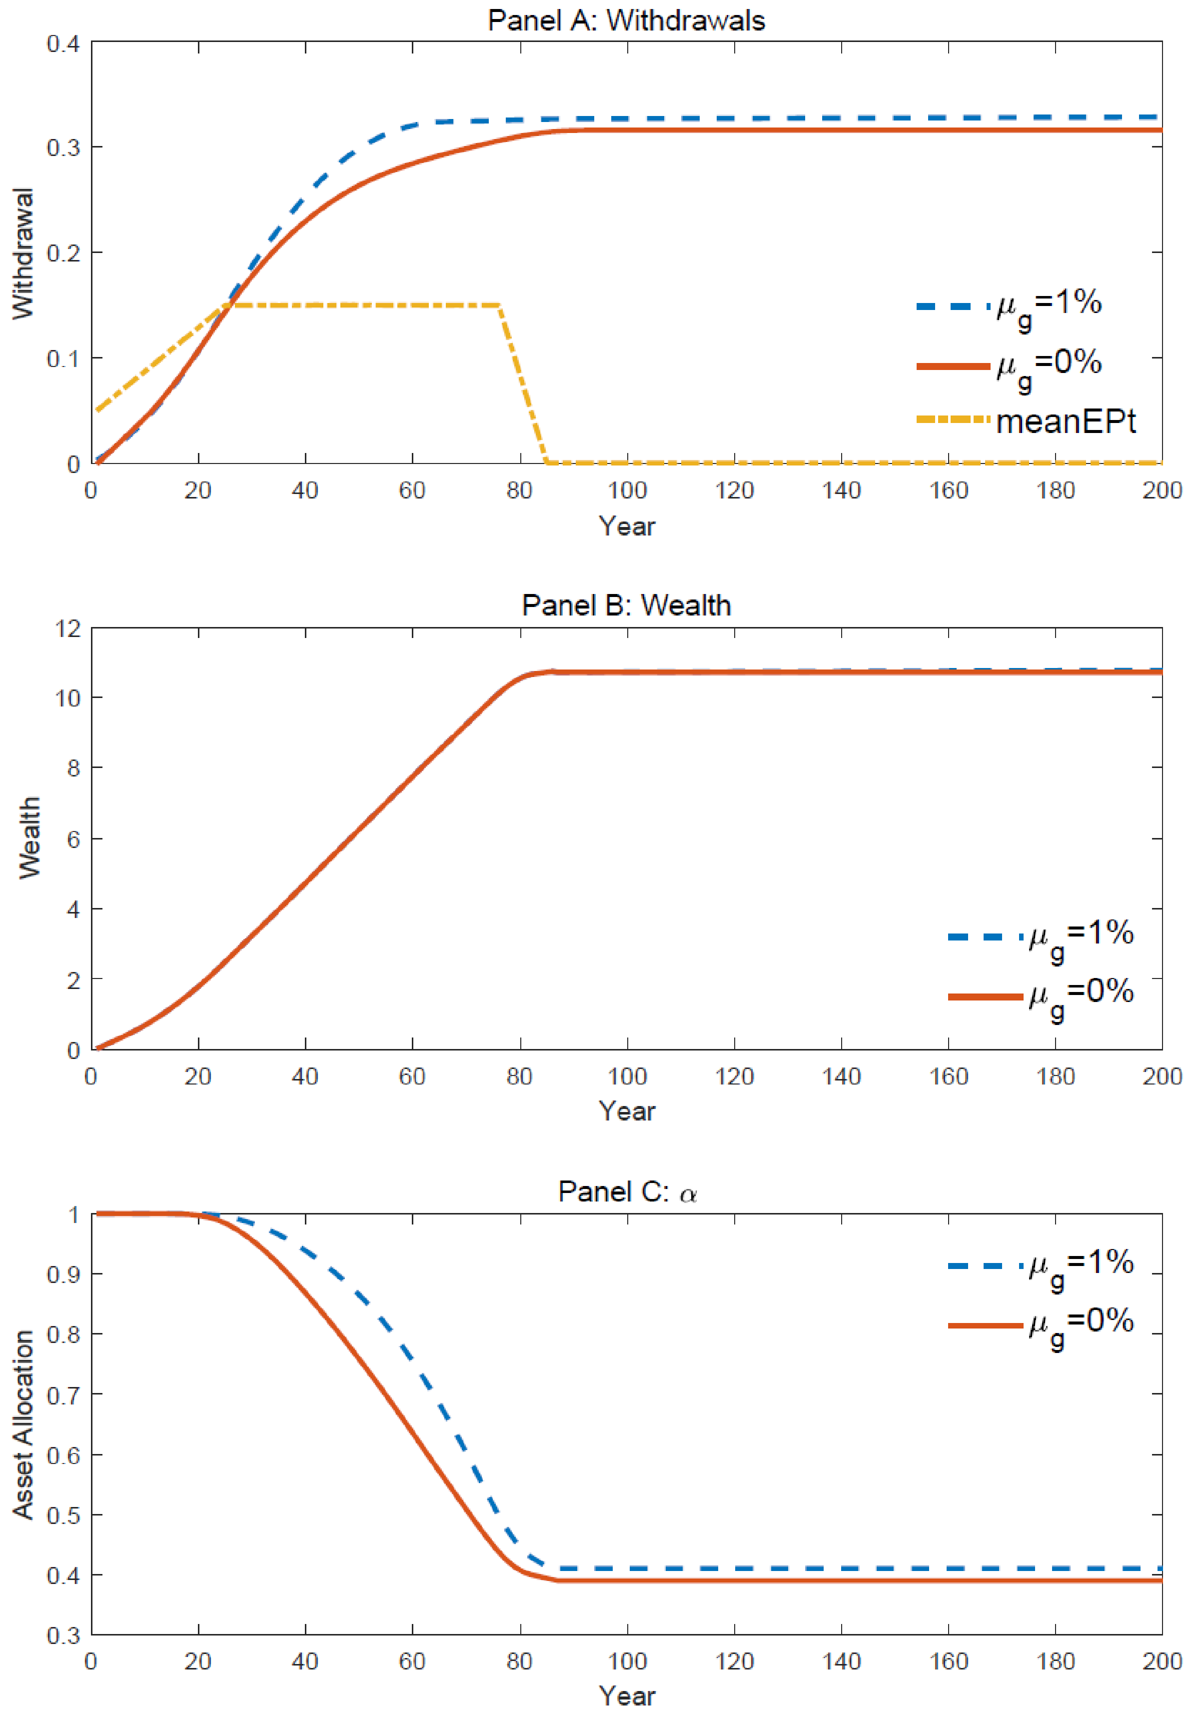

Figure 6B: Confidence Interval for model with $\mu_g = 0\%$

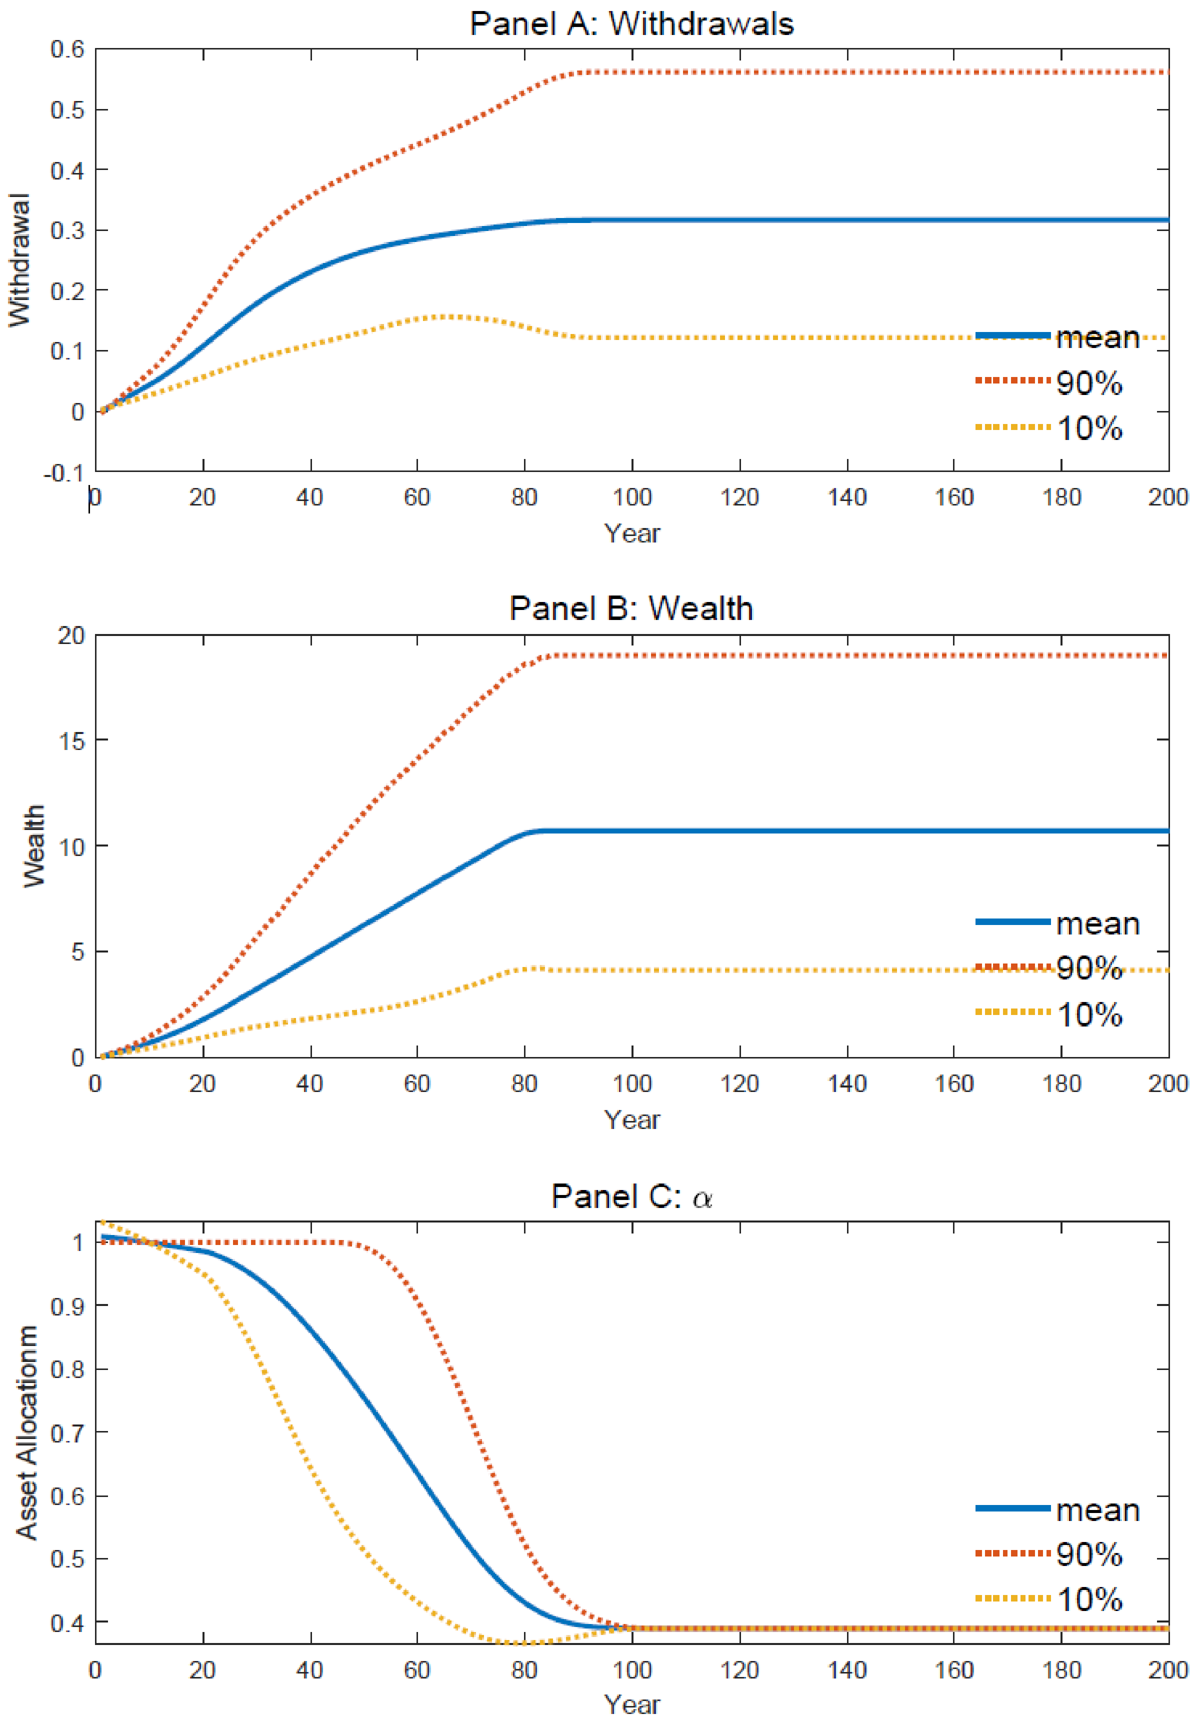

Figure 7A: Effect of Commodity Prices (ρ_{ye}=0.25)

Figure 7A presents the policy function from the baseline model where β=0.98, no habits, μ_g= 0%, and r_f=1% for low, mean and high commodity prices (f_t) and the mean GDP state (y_t). The correlation between the GDP and commodity price innovation (ρ_{ye}) is 0.25.

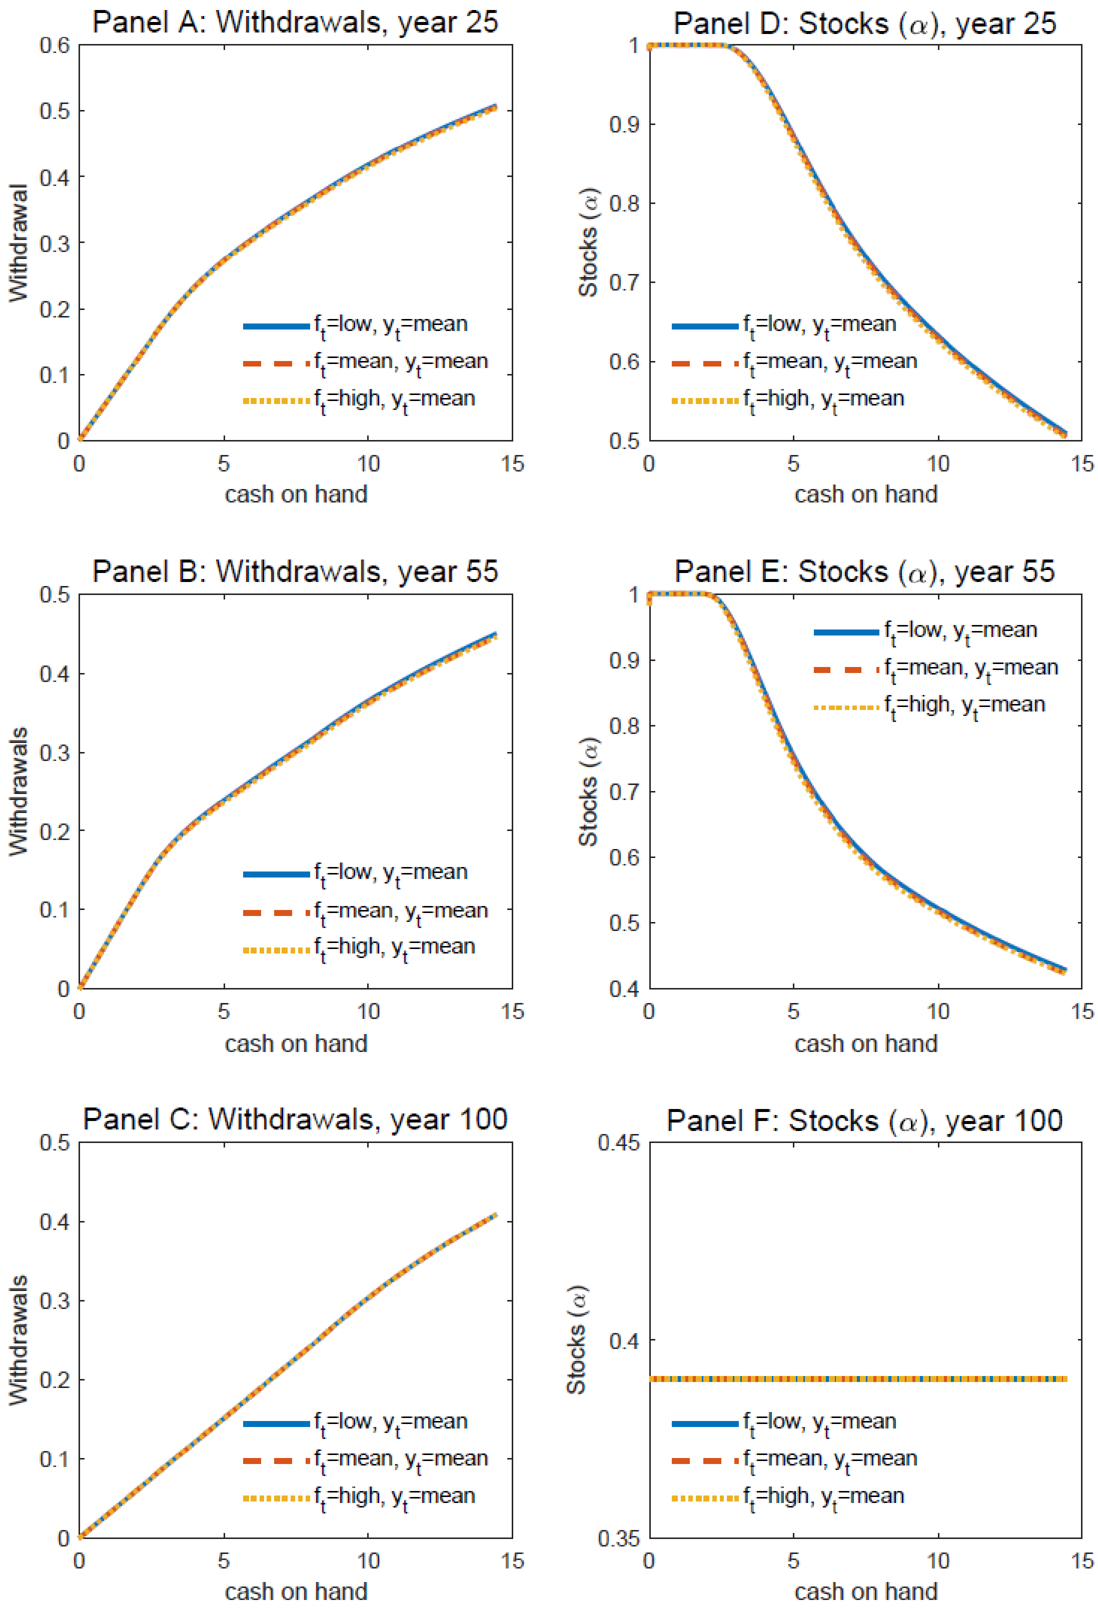

Figure 7B: Effect of Commodity Prices (ρ_{ye}=0.9)

Figure 7B presents the policy function from the baseline model where β=0.98, no habits, μ_g= 0%, and r_f=1% for low, mean and high commodity prices (f_t) and the mean GDP state (y_t). The correlation between the GDP and commodity price innovation (ρ_{ye}) is 0.9.

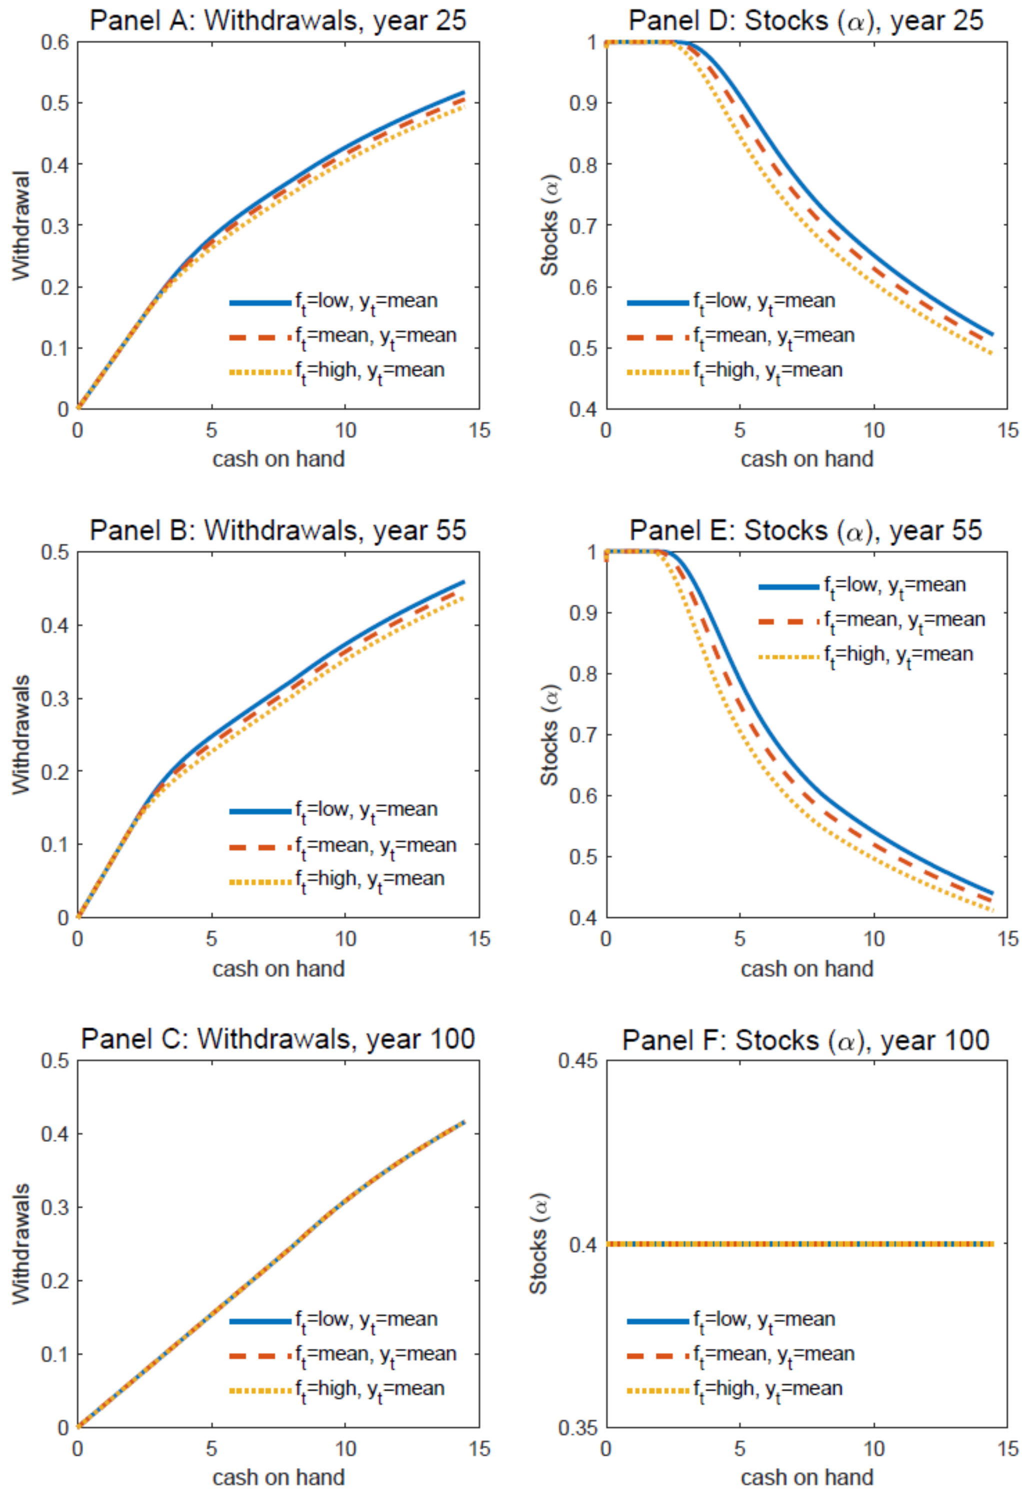

Figure 8A: Effect of the business cycle

Figure 8A presents the policy function results from the baseline model where β=0.98, no habits, μ_g= 0%, and r_f=1% for low, mean and high GDP (yt), and the mean commodity price state (ft). The correlation between the GDP and commodity price innovation (ρ_{ye}) is 0.25.

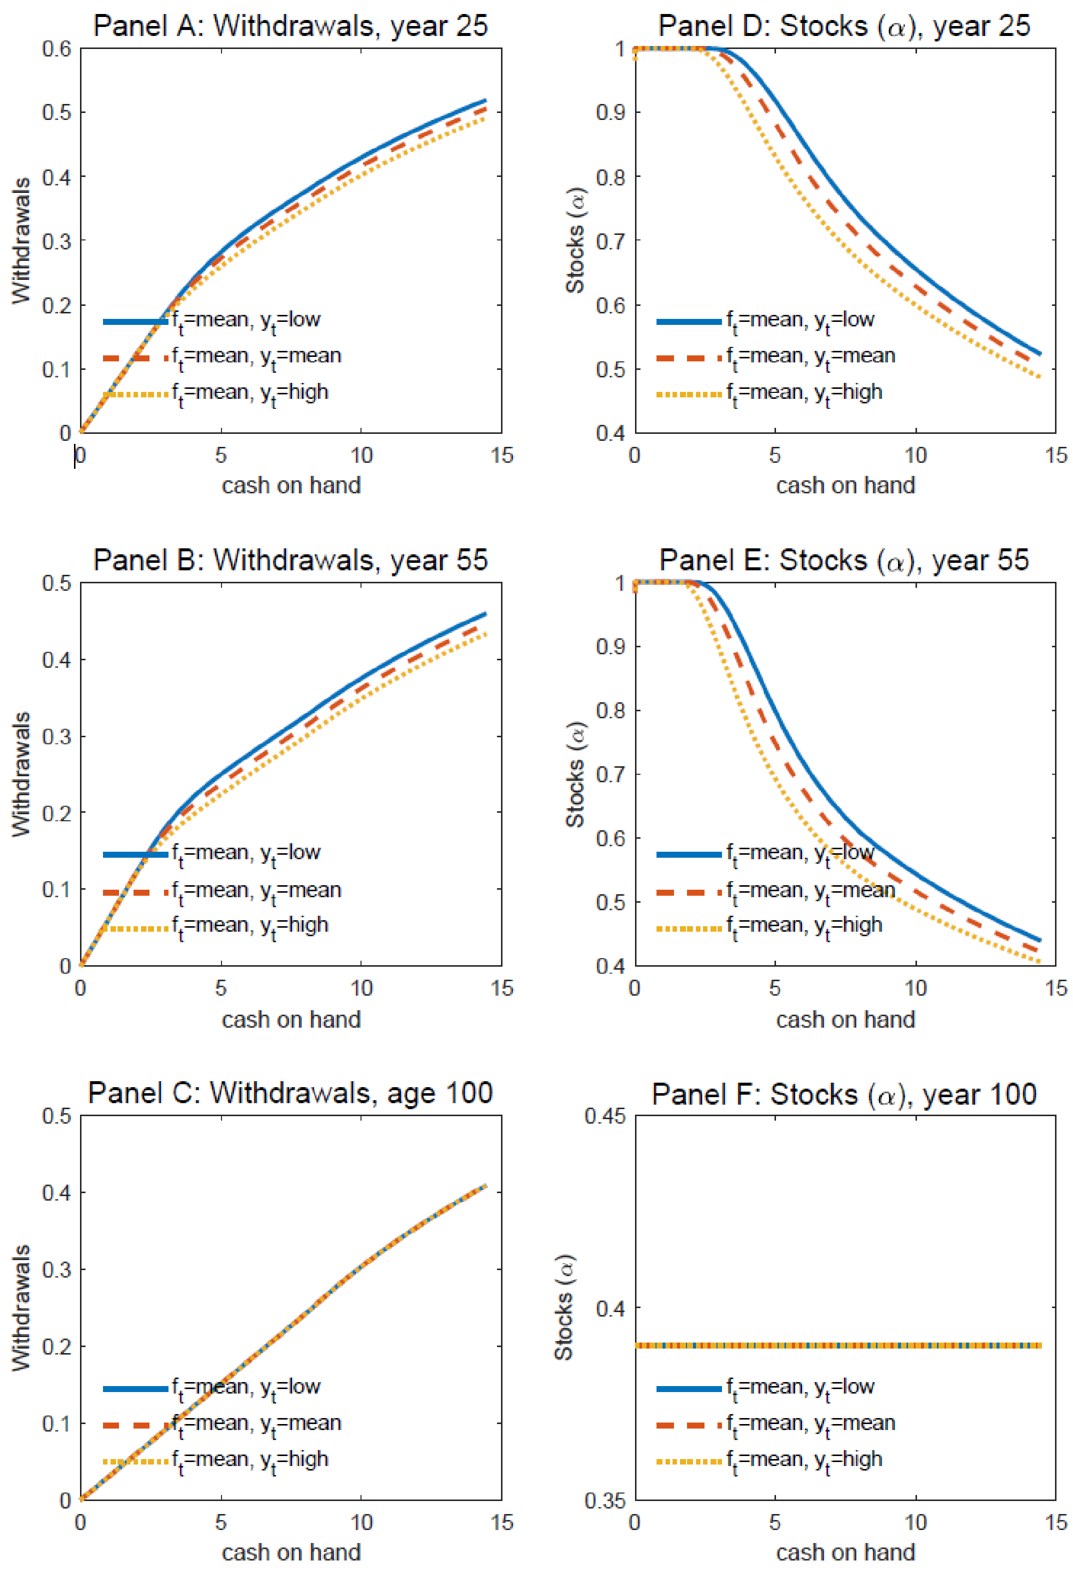

Figure 8B: Effect of the business cycle (ρ_{ve}=0.9)

Figure 8B presents the policy function results from the baseline model where β=0.98, no habits, μ_g= 0%, and r_f=1% for low, mean and high GDP (yt), and the mean commodity price state (ft). The correlation between the GDP and commodity price innovation (ρ_{ye}) is 0.9.

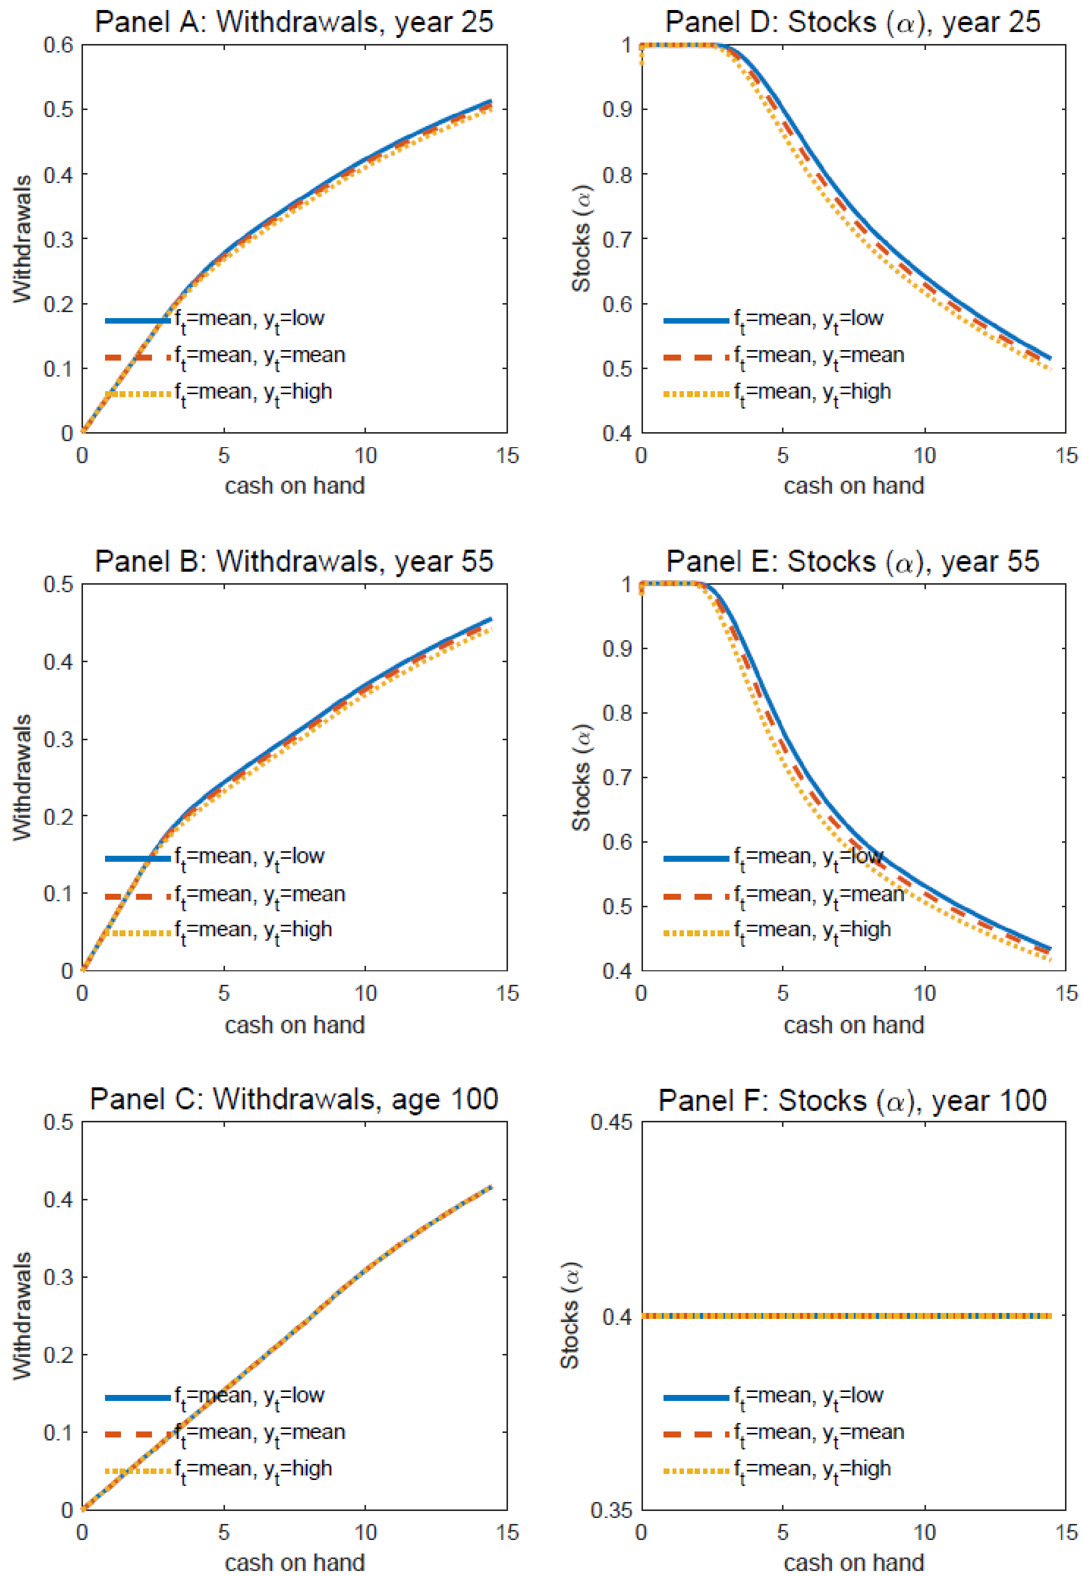

Figure 9: Simulations varying θ (taxation capacity)

Figure 9 presents the average results from the baseline model where θ=0.2 against the case when θ=0.0, ceteris paribus.

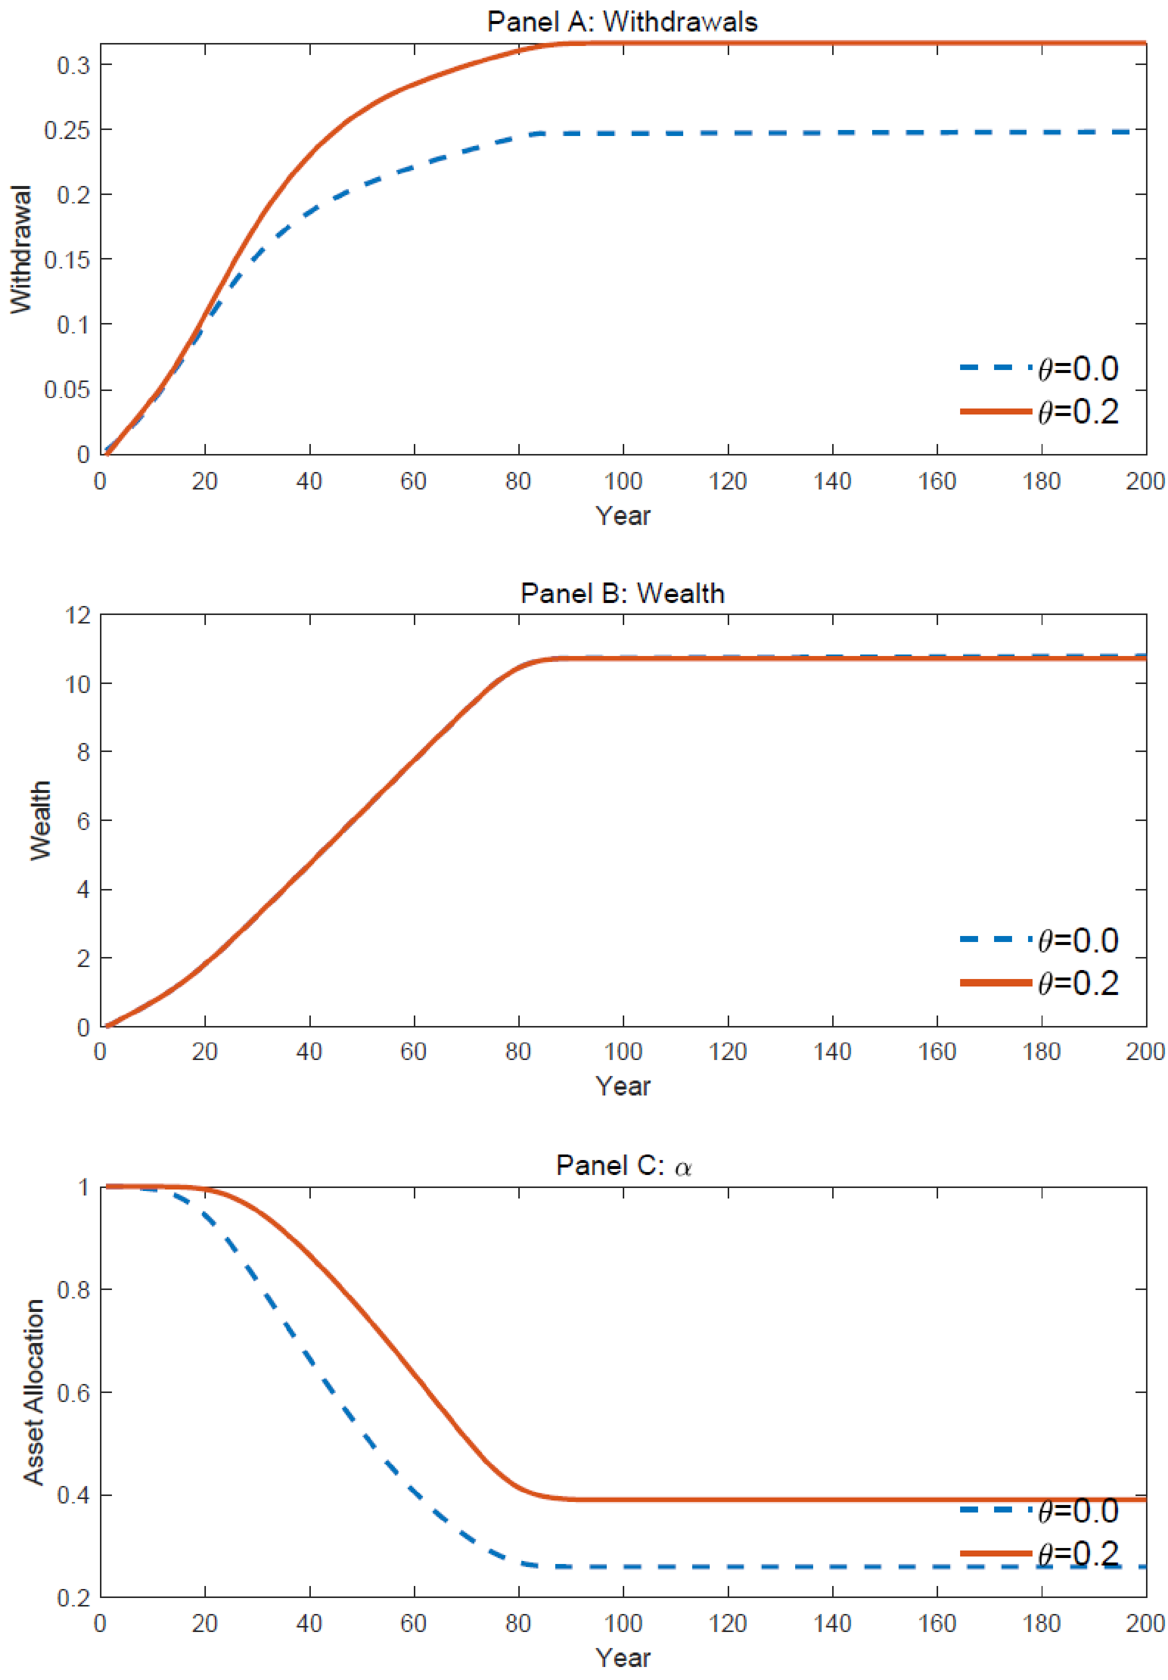

Figure 10: Effect of r_f

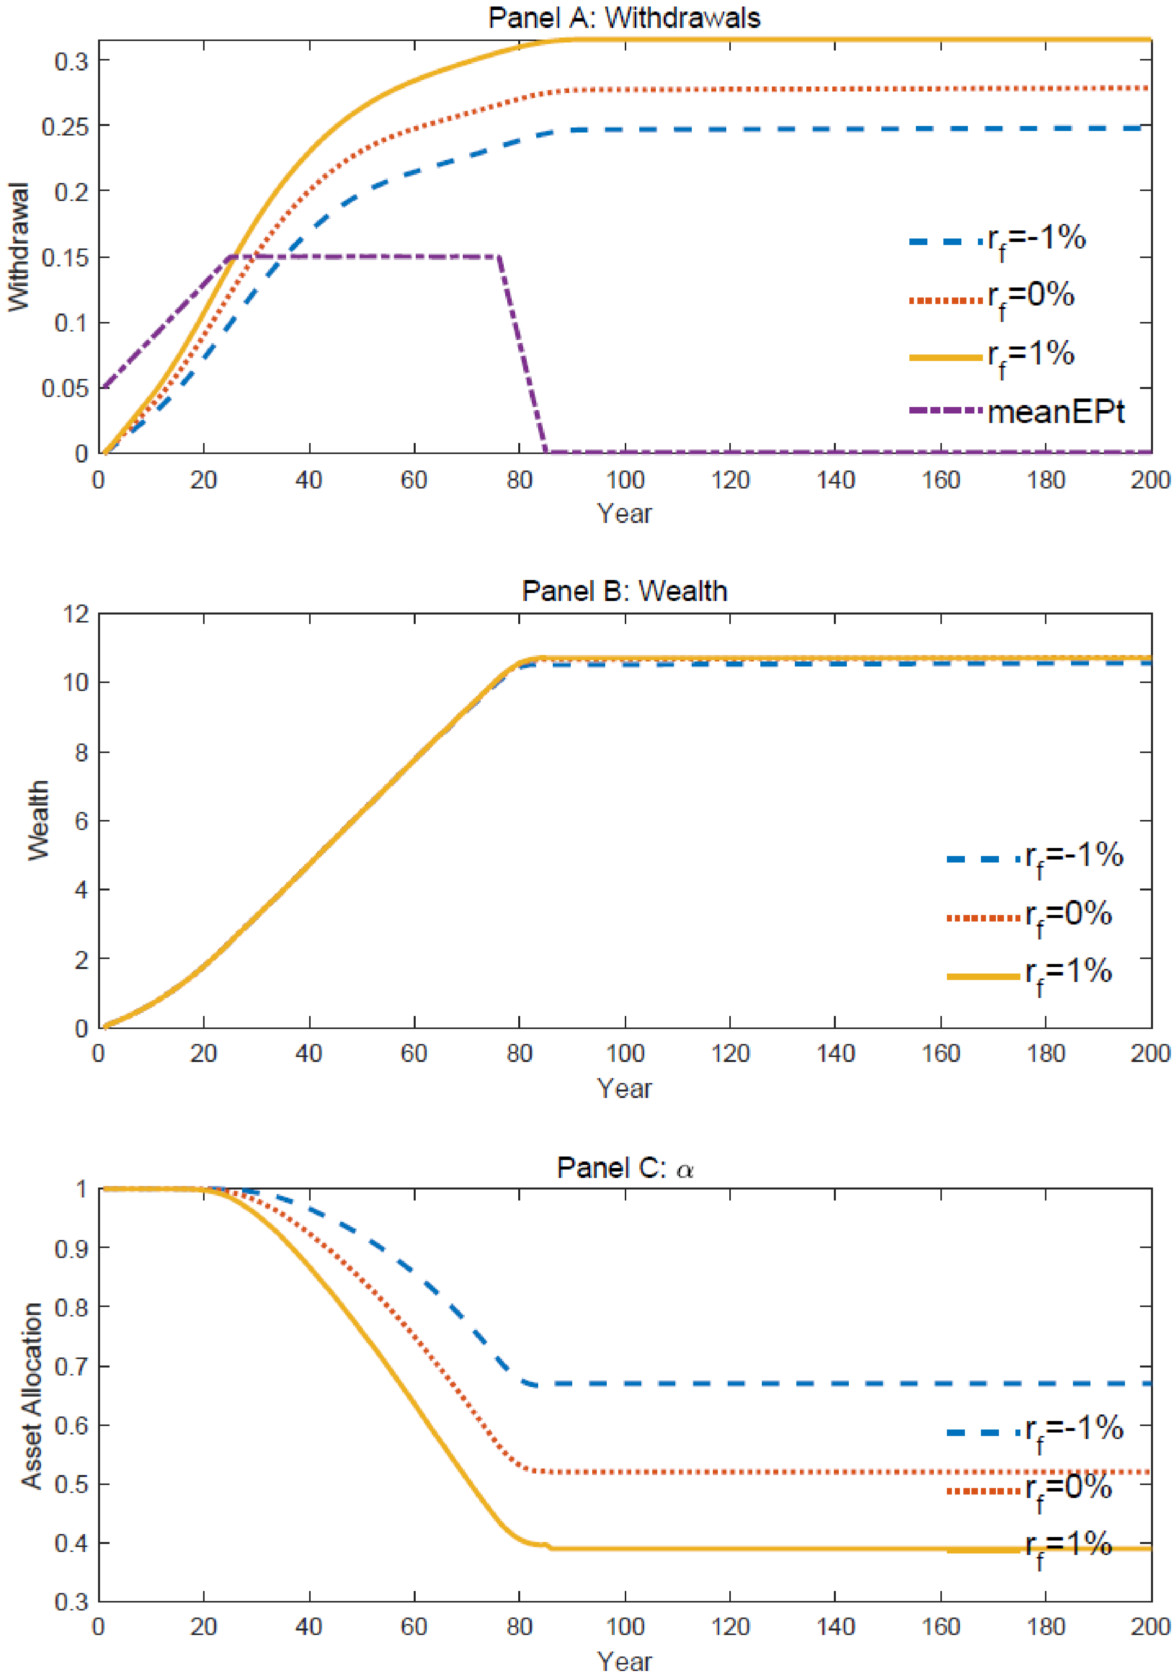

Figure 11: Capped Withdrawal Rate (Cap = 4%, Risk Premium = 5%, r_f=1%)

Baseline (μ_g=0%) relative to k-capped model where the withdrawal rate has to be less than 4% from the fund and relative to α-capped model where the share of wealth in stocks is less than 60%.

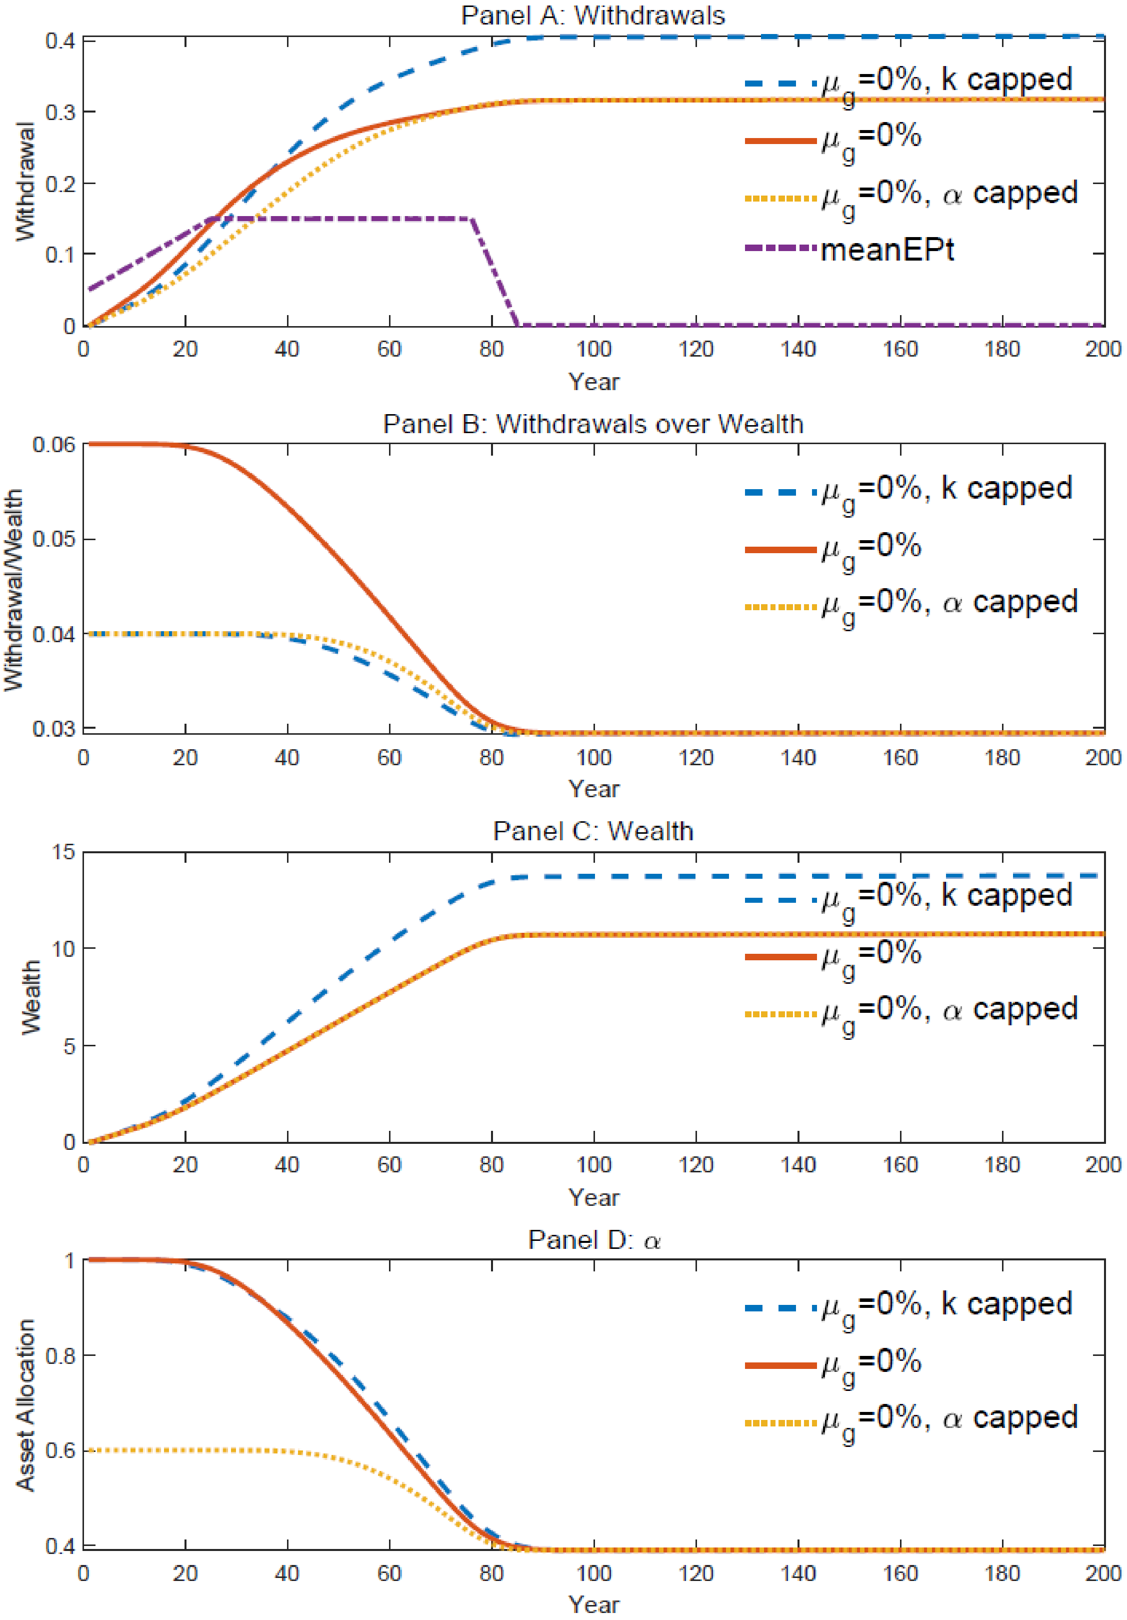

Figure 12: Effect of Habit (Risk Premium = 5%, r_f=1%)

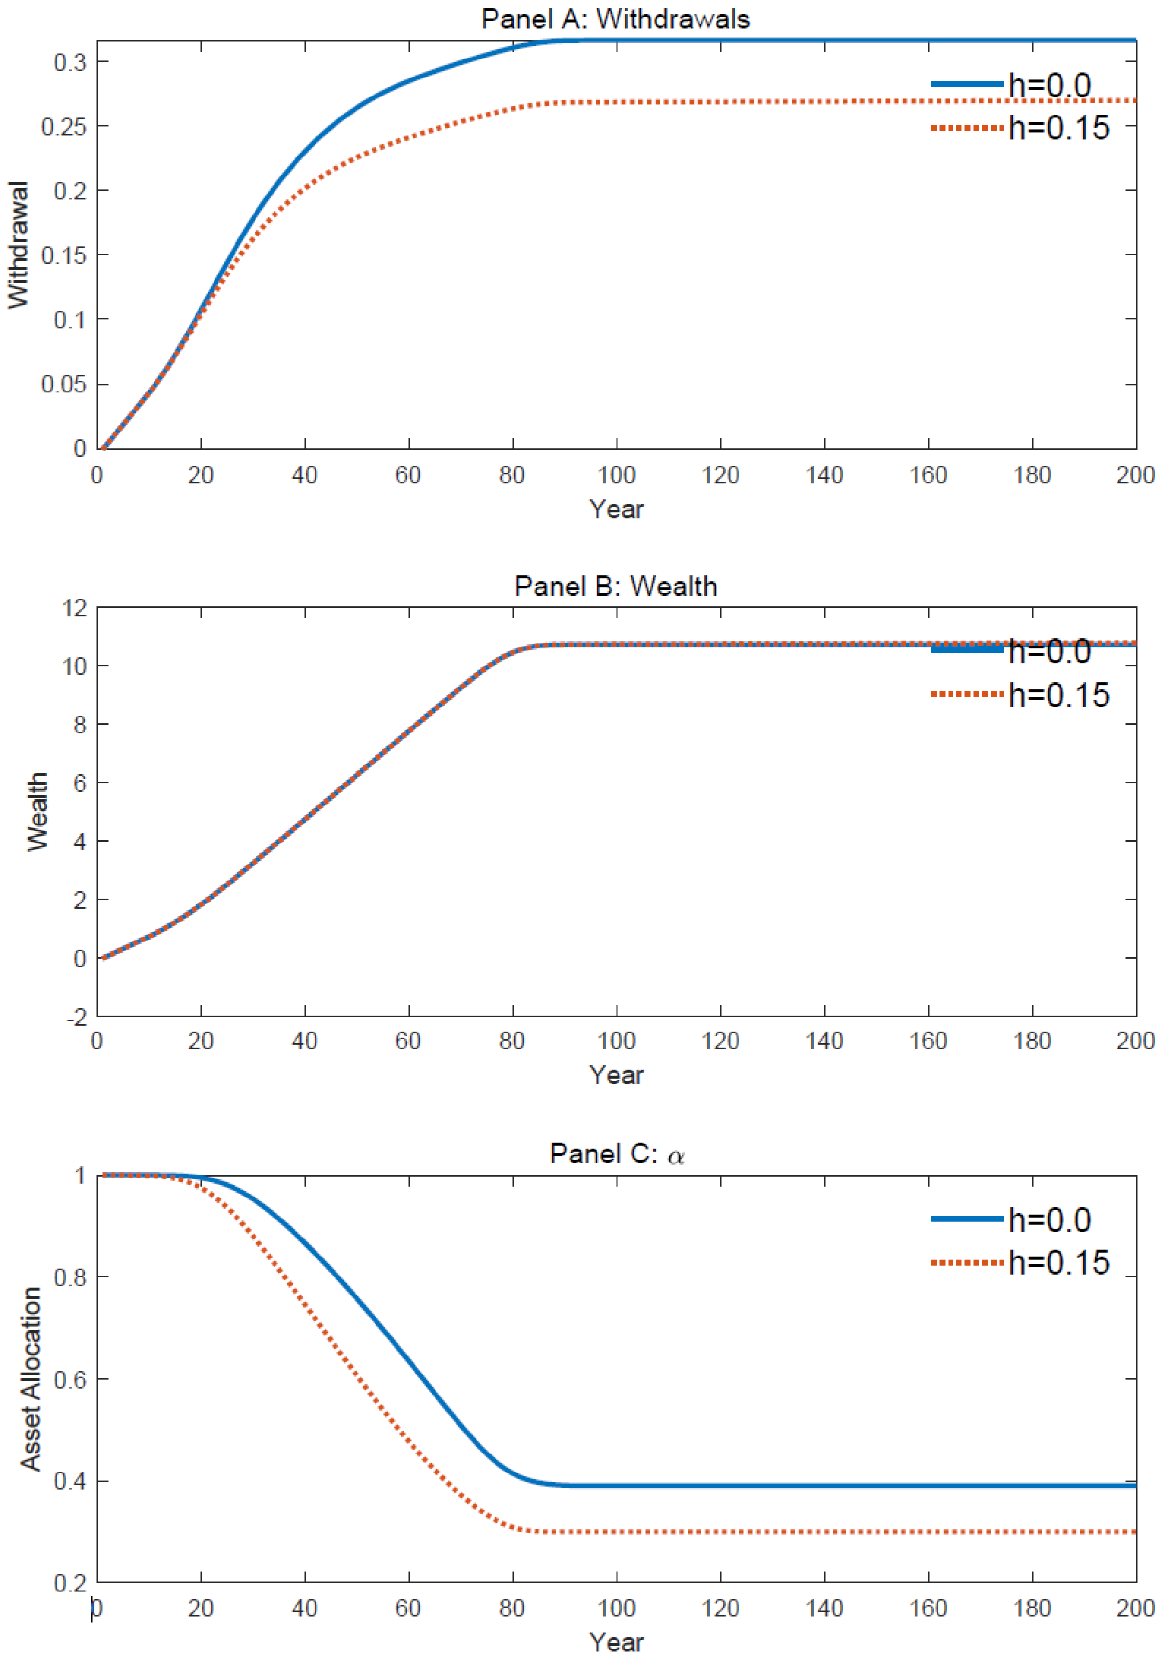

Figure 13: Effect of Different Model on α (Risk Premium = 5%, r_f=1%)

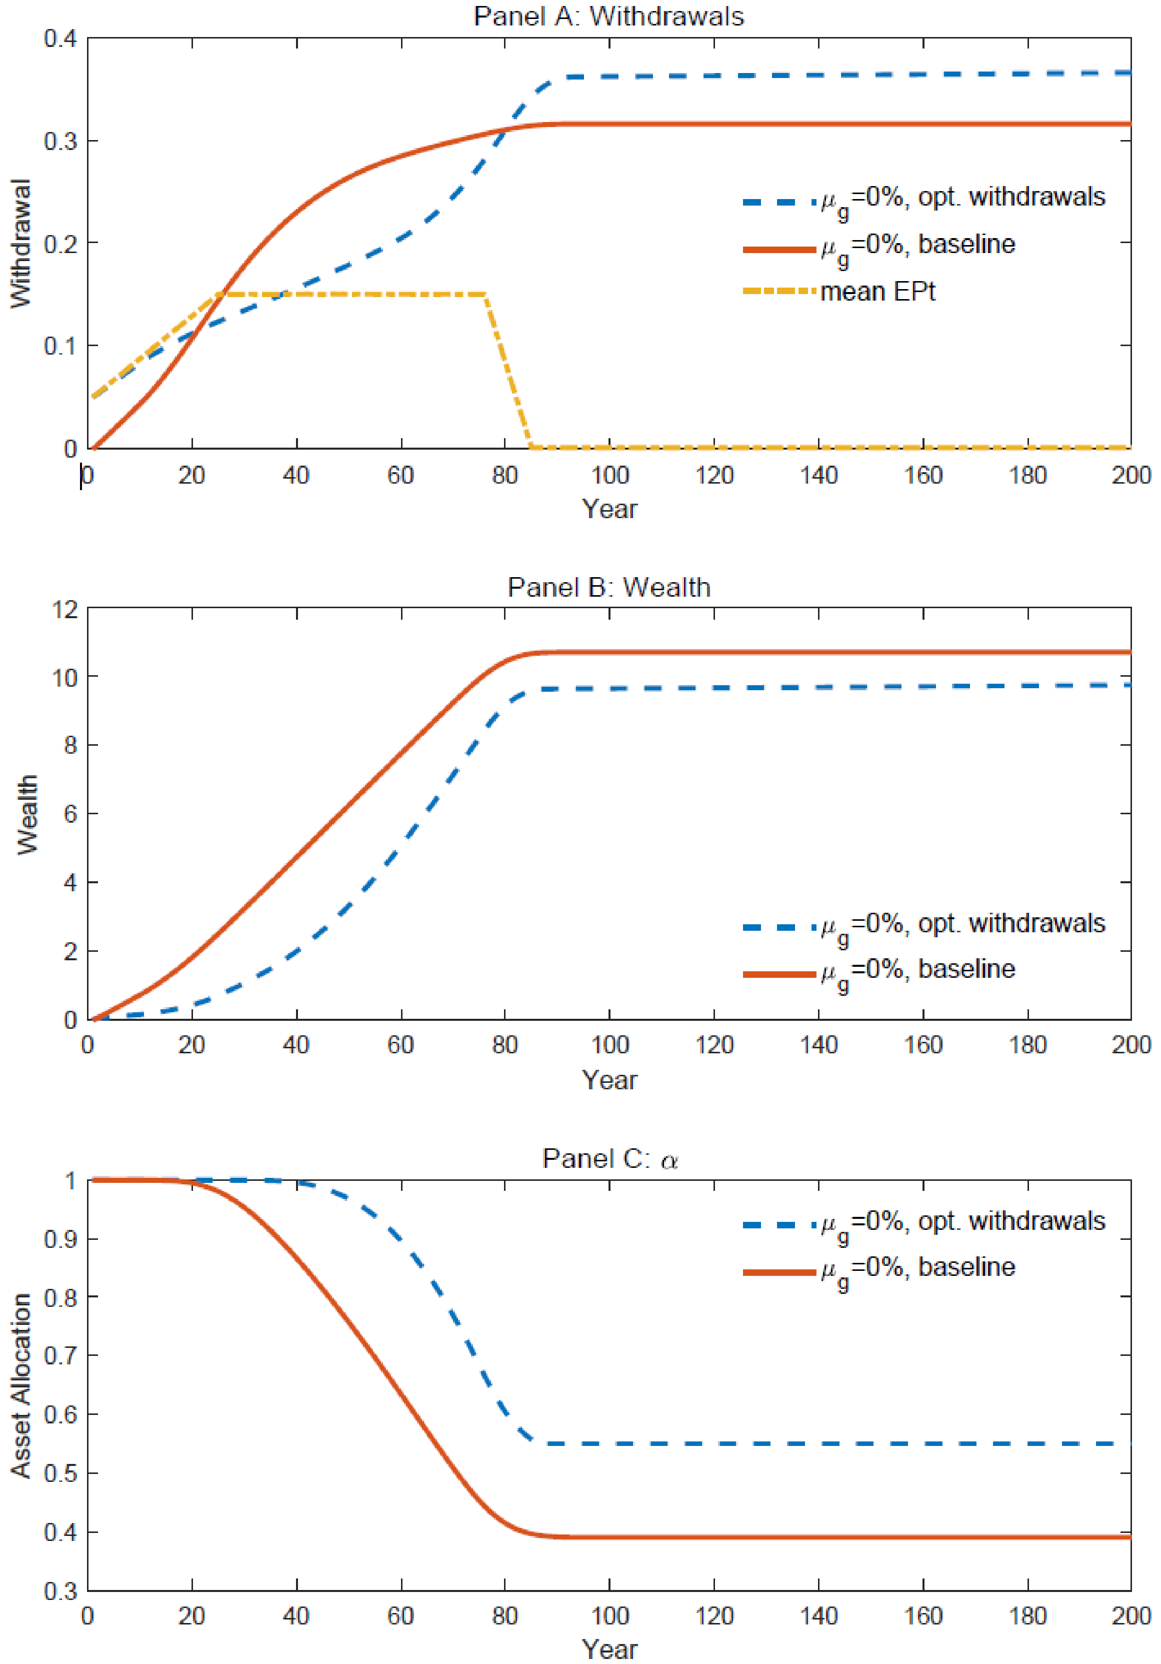

Figure 14: Effect of Different Model on α (Risk Premium = 5%, r_f=1%)