The Cost of ESG Rating Uncertainty*

Elise Gourier
ESSEC Business School and CEPR

Menglong Na
ESSEC Business School

Avril 2025

Abstract

More than thirty trillion dollars are allocated to ESG assets, despite the uncertainty on the true sustainability level of firms. We show within a dynamic portfolio allocation model that the effect of this uncertainty on investors' portfolios and welfare strongly depends on their preferences. Based on institutional investors' holdings, we identify two types of investors with non-pecuniary preferences: *green investors* whose utility increases with the estimated sustainability of firms, and *threshold investors*, whose non-pecuniary preferences only depend on whether a firm's estimated sustainability level is below or above a threshold. We show that the uncertainty of the ESG rating is twice as costly for green investors than for threshold investors.

JEL Classification: G11, G12

Keywords: ESG uncertainty, portfolio allocation, sustainable preferences

*Email: elise.gourier@essec.edu and menglong.na@essec.edu

1 Introduction

Over the past decade, asset managers have increasingly taken environmental, social and governance (ESG) standards into account in their investment decisions. ESG assets reached 30 trillion dollars of Assets Under Management (AUM) in 2022 and are expected to exceed 40 trillion dollars by 2030[1]. A major challenge for investors, however, is that there is no consensus on how to accurately evaluate firms' sustainability. Such evaluation faces two main challenges. First, ESG ratings differ in the criteria they evaluate, in the way each individual criterion is assessed, and in the aggregation of these criteria (Berg, Koelbel, and Rigobon (2022)). Second, increasing concerns have been raised in the past years about firms' misleading communications, a.k.a. greenwashing (Gourier and Mathurin (2024)). As a result, ESG ratings can be seen as noisy signals of firms' true sustainability levels. This paper quantifies the effect of this ESG rating uncertainty on investors' portfolio allocation decisions and welfare.

An empirical analysis of institutional investors' 13F filings reveals three types of investors, two of which tilt their portfolios toward green assets (Pastor et al. (2023)). These investors differ by size and activity. Small- to medium-size institutional investors invest less in greener assets, while large investors invest more. Among them, insurers invest more the greener the assets, whereas investment companies, banks and financial advisors appear to follow a threshold strategy whereby the allocation in greener assets stops increasing above a threshold ESG score. We label these three types of investors *brown investors*, *green investors*, and *threshold investors*. We model investments in ESG assets from the perspective of each type of investor, choosing preference specifications that are in line with their behaviors. Brown investors have a standard power utility. Green investors have standard non-pecuniary preferences, with utility that is monotonically increasing in the estimated sustainability of the portfolio. In contrast, threshold investors receive a constant level of additional utility if the portfolio's ESG rating is above a given threshold.

We solve a dynamic portfolio optimization problem for these three types of risk-averse investors with access to the market fund, bonds, and an ESG fund over an infinite horizon. The model takes the perspective of investors. They do not observe the fundamental sustainability level of the ESG fund, which takes positive values on a grid. Instead, they observe a stochastic signal, or ESG rating. The ESG rating is mean-reverting around the fundamental fund's sustainability level, to capture the persistent biases that greenwashing can introduce (Cartellier, Tankov, and Zerbib (2024)). A fund which invests in communication and tends

[1] https://www.bloomberg.com/professional/insights/sustainable-finance/esg-aum-set-to-top-40-trillion-by-2030-anchor-capital-markets/

to oversell its sustainability will naturally tend to have an ESG rating that is better than its fundamental level. In contrast, funds that invest less in public relations, or funds that suffered from a recent ESG incident, will likely be given a signal that underestimates their fundamental sustainability level. We distinguish between this type of ESG uncertainty, which we label *greenwashing uncertainty*, and the uncertainty that comes from the volatility of the stochastic signal. This volatility captures the measurement error associated with the ESG ratings (Avramov, Cheng, Lioui, and Tarelli (2022)). We refer to this second type of ESG rating uncertainty as *measurement uncertainty*. Our model is flexible enough to account for watchdogs, who, at times, report gaps between firms' perceived and fundamental sustainability level. These reports are modelled by a Poisson process, with intensity increasing in the size of the gap. When this process jumps, the gap is brought to zero.

To model the changes in ESG funds' strategy, or more generally in the strategy of their underlying firms, the ESG fund's fundamental sustainability level is a jump process. It switches value on the grid at random times, modeled with a Poisson process. Neither these switches in value nor their time of occurrence are observed by investors. Investors thus estimate the true sustainability level of the ESG fund as the value on the grid that is the closest to the signal.

In line with Pastor et al. (2022), expected returns on the ESG fund are lower than those on the market fund. The more sustainable the ESG fund, the lower its return. Investors therefore face a trade-off between higher expected returns (pecuniary utility) and higher sustainability (non-pecuniary utility). They choose their optimal portfolio allocation and consumption policy based on the estimated fund sustainability and corresponding returns. The portfolio is continuously rebalanced.

We calibrate our model to data and reproduce remarkably well the patterns in the green allocations of the three types of investors. In line with the data, CRRA investors have an allocation to the ESG fund that decreases with the ESG signal. In contrast, green investors have an allocation that increases with the ESG signal before the signal reaches the threshold, but remains flat afterwards. We use the model to address two questions: Do greenwashing uncertainty and measurement uncertainty matter to investors? Do they matter differently to the three types of investors?

We answer these questions using two measures, following Gourier, Phalippou, and Westerfield (2024). First, the welfare cost or gain is defined as the one-off percentage of wealth investors would give up, or pay, to remove greenwashing (respectively, measurement) uncertainty. Removing greenwashing uncertainty amounts to setting the drift of the signal equal to the ESG fund's fundamental sustainability level, whereas removing measurement

uncertainty sets the volatility of the signal equal to zero. Second, the return premium is the permanent return loss or gain on the ESG fund that investors would accept to remove the uncertainty.

Our results are twofold. First, greenwashing uncertainty matters, whereas measurement uncertainty does not. Specifically, greenwashing uncertainty affects the welfare of investors with non-pecuniary preferences, i.e., green and threshold investors, whereas measurement uncertainty carries a negligible cost to all types of investors. Second, greenwashing uncertainty affects green and threshold investors differently. With our calibrated parameters, the welfare cost of ESG level uncertainty is positive. The lower the ESG signal, the higher the welfare cost, which reaches 4% of wealth at its maximum. In contrast, for threshold investors, the welfare cost reaches a maximum of 2%. That is, both green and threshold investors prefer certainty over uncertainty of the ESG fund's sustainability level; threshold investors however have less incentives spending money to improve the signal. The return premium behaves in a similar way.

We use a sensitivity analysis to study the impact of parameters on the sign of the welfare cost and return premium. Decreasing the intensity of random reports of greenwashing decreases the welfare cost of ESG rating uncertainty. Removing these reports completely even makes investors prefer uncertainty over certainty of the ratings, for high values of the signal. This result underlines the importance of having watchdogs report greenwashing, to incentivize institutional investors to improve the quality of the signal.

The rest of the paper is organised as follows. Section 2 reviews the literature we are building on. Section 3 studies the portfolio allocations of institutional investors and highlights three types of investors. Section 4 presents our model and Section 5 describes how it is calibrated. Section 6.1 presents the optimal portfolio allocations of the different types of investors, and describes the welfare costs associated with greenwashing and measurement uncertainty.

2 Literature Review

We build on two main strands of literature. The first strand highlights the uncertainty of ESG ratings, and the mistrust of investors in these ratings. Krueger, Brandon, and Schmidt (2021); Berg, Koelbel, and Rigobon (2022); Christensen, Serafeim, and Sikochi (2022) and Berg, Koelbel, Pavlova, and Rigobon (2023) find that there is considerable disagreement across rating agencies regarding how to view a given firm's ESG performance, resulting in low pairwise correlations between the different ratings of a given firm. Using textual

analysis, Gourier and Mathurin (2024) show that ESG ratings have been described in the media as sources of confusion and mistrust for investors. Cartellier, Tankov, and Zerbib (2024) show within a general equilibrium model that firms had incentives to greenwash, in order to introduce upward biases in their ratings. We decompose the total uncertainty of ESG ratings into measurement uncertainty and greenwashing uncertainty. We model each source of uncertainty within a partial equilibrium setup, and assess their effect on investors' welfare. Avramov, Cheng, Lioui, and Tarelli (2022) and Krueger, Brandon, and Schmidt (2021) suggest that measurement uncertainty leads to a higher risk premium. We find that greenwashing uncertainty carries a larger cost for investors.

We also build on a second strand of literature, which studies the role of green assets in portfolio allocation. As the ESG level of firms is not directly observable, the effects of a firm's sustainability on its returns have been the subject of several papers, although there has not been a clear consensus. Results largely depend on the sustainability measure that is used, but also on the time horizon over which returns are calculated, and on the asset considered.

Many papers have found that sustainable firms have lower returns. Hong and Kacperczyk (2009) and Chava (2014) find that sin stocks outperform other stocks. Recent studies have found supporting evidence for this effect, e.g., Bolton and Kacperczyk (2021), Bolton and Kacperczyk (2023) and Hsu, Li, and Tsou (2022) find that companies emitting the most greenhouse gases earn higher stock returns than companies emitting the lowest levels. A positive return premium on brown investments in equities, bonds, real estate, and option markets is also documented in Faccini, Matin, and Skiadopoulos (2023); Huynh and Xia (2021); Seltzer and Zhu (2022); Bernstein, Gustafson, and Lewis (2019); Baldauf, Garlappi, and Yannelis (2020); Painter (2020); Goldsmith-Pinkham, Gustafson, Lewis, and Schwert (2023) and Ilhan, Sautner, and Vilkov (2021).

Several explanations have been put forward to justify lower returns of sustainable firms. Green investors, along the lines of Fama and French (2007), may push up asset prices of green firms and reduce their cost of capital, Heinkel et al. (2001). They may do so because of non-pecuniary preferences, Fama and French (2007); Friedman and Heinle (2016); Pastor, Stambaugh, and Taylor (2021); Pedersen, Fitzgibbons, and Pomorski (2021a); Avramov, Cheng, Lioui, and Tarelli (2022); Landier and Lovo (2023), because they internalize the cost of future regulations or climate change externalities, Zerbib (2022); De Angelis, Tankov, and Zerbib (2023), or because of impact mandates they have, Oehmke and Opp (2024). A lower cost of capital allows sustainable firms to implement costly ESG actions, and become a hedge against adverse physical and transition risks.

Conversely, several papers find that sustainable stocks have higher returns than other stocks, Edmans (2011); Khan et al. (2016); Lins et al. (2017); Hong et al. (2019); ?); Görgen et al. (2020); Berg et al. (2024) . Several explanations for the positive return premium of sustainable stocks have been proposed. The first explanation is under-reaction to ESG information, Glossner (2021). Pedersen, Fitzgibbons, and Pomorski (2021a) propose an equilibrium model with green preferences and ESG scores that are informative about stocks' risk and return. Their model shows that green stocks can have higher returns if ESG scores positively predict returns in a way that has not been appreciated by all investors. Second, sustainable assets may benefit from increased product differentiation and higher profit margins, Albuquerque et al. (2019). A third explanation is demand pressure. Pedersen, Fitzgibbons, and Pomorski (2021b) and van der Beck (2022) document persistent demand pressure that pushes green stocks' prices up. van der Beck (2022) finds that without this flow-driven effect, sustainable funds would have underperformed in the period under consideration. An important common denominator of all these mechanisms is that the outperformance of green firms should be temporary. Leippold and Yu (2023) and Avramov, Cheng, Lioui, and Tarelli (2022) find empirical evidence for such temporary effects. Finally, Pastor, Stambaugh, and Taylor (2022) highlight the distinction between ex-ante expected returns and ex-post realized returns. Whereas U.S. green stocks outperformed brown stocks during the last decade, they find that expected returns on green stocks are on average 1.4% per year lower than those on brown stocks.

The horizon over which returns are calculated is therefore crucial, in line with market prices taking years to fully reflect the value of intangibles. Firms with high R&D (Lev and Sougiannis (1996), Chan et al. (2001)), advertising (Chan et al. (2001)), patent citations (Deng et al. (1999)), and software development costs (Aboody and Lev (1998)) all earn superior long-run returns. The outperformance of firms with high employee satisfaction that is shown by Edmans (2011) is a long-term result. ESG industry proponents widely share the view that sustainable firms outperform on the long-term, see, e.g., Times (2017), BlackRock (2020); Paribas (2019), and Schroders (2020).

Our model allows to study the trade-off between investors' non-pecuniary preferences, and the loss in expected returns of green assets. We introduce a new type of non-pecuniary preferences, whereby utility does not increase with the ESG signal when the signal is above a given threshold. We present novel empirical evidence for investors with such preferences, and examine the impact of ESG rating uncertainty on the welfare of these investors. We also examine an extension that accounts for long-term outperformance of green assets, but this extension does not change our results.

3 Green Tilts of Institutional Investors

In this section, we analyze the holdings of institutional investors, to better understand their preferences in terms of green assets. We build on the methodology of Pastor et al. (2023) and compute the tilt of groups of investors for assets with varying sustainability levels. We identify three types of investors. Brown investors have a tilt function that is decreasing in the level of sustainability of their holdings. Green investors, in contrast, have a tilt function that increases in sustainability. Threshold investors' tilt function is non-linear in holdings' sustainability level: it is increasing up to a threshold, and flat above this threshold.

3.1 Data

We use the MSCI ESG ratings of stocks, as of 2021, to assess the sustainability of companies. MSCI ESG ratings range from 0 to 10. Ratings are relative, i.e., companies are assessed in comparison of their peers. Ratings are available for 1987 U.S. stocks.

The holdings of institutional investors are obtained from the 13F filings that the SEC requires investment companies and professional money managers to file, as of December 2021.[2] We filter out investors who hold less than 50 stocks with available ESG rating and firm characteristics. After this step, 2733 filings are available, for total Assets Under Management (AUM) of 21.8 trillion US dollars. Institutional investors are divided into five categories, as described in Table 1: 1548 investment advisors (20% of total AUM, essentially small- and medium-sized),[3], banks (14% of total AUM, mostly large), insurers (2% of total AUM, almost all large), investment companies (1% of total AUM, mostly small), and other (62% of total AUM, most of all large). The "other" category includes pension funds and university endowments.[4]

We obtain firm characteristics from the CRSP and Compustat dataset. We extract the firm's share price, dividends per share, book value per share, number of shares outstanding, and market capitalization at the end of 2021.

[2]Investors do not have to report short positions, hence our data does not contain any information on these short positions.

[3]Assets under management (AUM) are estimated from the reported holdings of stocks with an ESG rating. Small caps and micro-caps are therefore not included in the calculation.

[4]Similar to Pastor et al. (2023), we ignore the holdings of stocks that do not have an ESG rating, so investors' total assets under management are calculated as the sum of all stocks with an ESG rating. The holdings of the stocks we cover account for 72% of the total holdings of institutional investors.

3.2 Green tilts of institutional investors

We calculate the green, brown and Green Minus Brown (GMB) tilts of each type of institutional investor following the methodology of Pastor et al. (2023). The green (respectively, brown) tilt of a given type of investors captures how much this type of investors allocate to green assets, relative to the market's holdings. The green tilt increases if investors increase their allocation to green assets, relative to the market, and/or if they decrease their allocation to brown assets. Stocks are considered green or brown depending on their ESG score. If the ESG score g_n of stock n is above the median ESG score across all stocks, g_0, then the stock is green. If g_n is below g_0, it is brown.

Let us consider investor i. For each stock n held by this investor, we calculate the tilt of investor i with respect to that stock, Δ_{in}. This tilt is defined as the expected allocation of investor i to stock n relative to the market, given a set of characteristics of the stock:

$$\Delta_{in} = \mathbb{E}\left(w_{in}|\mathcal{G}, \mathcal{C}\right) - \mathbb{E}\left(w_{in}|\mathcal{G}_0, \mathcal{C}\right) \tag{1}$$

The weight of stock n in the portfolio of investor i is denoted by w_{in}. $\mathcal{G}$ contains the ESG characteristics of all stocks, i.e., in our application, their ESG rating. $\mathcal{G}_0$ is the counterpart of $\mathcal{G}$ for the market portfolio. $\mathcal{C}$ contains firms' non-ESG characteristics: firm size, return of the past 12 months, book to market ratio, dividend rate and profitability. Profitability is measured as the gross profit divided by the book value of the firm. We estimate Δ_{in} using OLS regressions as described in Pastor et al. (2023).

The tilts with respect to all stocks in investor i's portfolio are then aggregated into a total tilt, by distinguishing between over-allocation: $\Delta_{in}^{+} := \Delta_{in} 1_{\Delta_{in}>0}$, and under-allocation: $\Delta_{in}^{-} := -\Delta_{in} 1_{\Delta_{in}<0}$, relative to the market. The green tilt of an investor, T_i^G, comes from their over-allocation to green stocks, $\Delta_{in}^{+} 1_{g_n>g_0}$, and their under-allocation to brown stocks, $\Delta_{in}^{-} 1_{g_n \leq g_0}$:

$$T_i^G = \sum_{n=1}^{N} \Delta_{in}^{+} 1_{g_n>g_0} + \sum_{n=1}^{N} \Delta_{in}^{-} 1_{g_n \leq g_0}. \tag{2}$$

Figure 1, Panel A, shows, in blue, the weighted average green tilts of institutional investors in each category covered by the 13F form. The weights are proportional to the assets under management of each investor in the category. Insurers have the largest green tilt, followed by investors in the category "other", banks, advisors and investment companies.

The brown tilt of investor i can be similarly defined as coming from the over-allocation

to brown stocks, and from the under-allocation to green stocks:

$$T_i^B = \sum_{n=1}^{N} \Delta_{in}^{+} 1_{g_n \leq g_0} + \sum_{n=1}^{N} \Delta_{in}^{-} 1_{g_n > g_0} \tag{3}$$

Figure 1, Panel A, shows, in red, that advisors and investment companies have the largest brown tilts, and banks and insurers the lowest. This ranking is close to the ranking by green tilt. The magnitude of brown tilts is small: institutional investors, overall, do not over-weigh brown stocks more than the market, neither to they under-weigh green stocks.

Finally, the green-minus-brown (GMB) tilt is simply the difference between the green and the brown tilts of investor i:

$$T_i^{GMB} = T_i^G - T_i^B \tag{4}$$

As green tilts are much larger, in magnitude, than brown tilts, the GMB tilts, in yellow, are very similar to the green tilts. These results are in line with the ones reported in Pastor et al. (2023).

Also consistent with the results of Pastor et al. (2023), we find that larger institutional investors have larger green tilts than smaller investors (even though the calculation of green tilts controls for firm size). In Figure 1, Panel B, we sort investors by assets under management, split them in three terciles, and display their tilts. The total AUM of small, medium and large investors accounts for $0.3\%, 1.2\%, 98.5\%$ of the total AUM, respectively. The green tilt increases with size, the brown tilt decreases, and as a result the GMB tilt increases sharply.

We conclude that the greenness of institutional investors varies across types and sizes of investors.

3.3 Green tilts as a function of asset greenness

The green tilts that we have computed split assets in two buckets: green and brown assets, based on their ESG score relative to g_0 –the median ESG score–. We now separate assets in 40 buckets, to distinguish between assets whose score is close to g_0, and assets whose score is far above or below g_0. The buckets are formed by sorting stocks by their ESG score, such that all buckets contain the same number of stocks.

We now define investor i's tilt for a given bucket of greenness G_j, $j \in \{1, ..., 40\}$ as:

$$T_i^{G_j} = \sum_{n=1}^{N} \Delta_{in} 1_{g_n \in G_j} \tag{5}$$

Figure 2, Panel A, displays the weighted average tilts $T_i^{G_j}$ of each type of institutional investors, as a function of the lower bound of each ESG bucket. This graph uncovers two types of investors. First, insurers have the lowest tilt for brown stocks (left side of the graph) and the highest tilt for green stocks. Their tilt increases monotonically with stocks' ESG ratings. Second, the tilt of banks, investment companies and advisors is almost flat before and above the median g_0, at which there is an upward jump.

Figure 2, Panel B displays the weighted average tilts $T_i^{G_j}$ grouping investors by size, and uncovers a third type of investors. Small and medium-size investors have a tilt that decreases when the stocks' ESG rating increases. In contrast, the tilt of large investors follows the pattern described in Panel A and increases with the stocks' ESG rating.

Thus, this analysis uncovers three types of investors. The first type, insurers, tilt their portfolio more towards greener assets, in line with them having non-pecuniary preferences increasing with their portfolio's greenness level. The second type (banks, investment companies and advisors) behave like threshold investors: they tilt their portfolio towards green assets until they reach a given level of greenness. Beyond this level of greenness, their tilt remains flat. The last type (small and medium-size investors) tilts their portfolio towards brown assets. Assuming that it is more costly to invest in green assets, this behavior is in line with standard preferences with no non-pecuniary component.

4 Model

We model investment portfolios that combine the market index, an ESG fund and a riskless asset. Our setup is designed to capture the preferences of the three types of institutional investors uncovered in Section 3.

4.1 Economy

There are three liquid assets in the economy that can be rebalanced continuously at no cost: a risk-free bond, which captures the fixed income market, a market index, which captures

the public equity market, and an ESG fund.[5]

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

$$dB_t = rB_t dt, \tag{6}$$

and the stock price P_t grows as a Geometric Brownian Motion:

$$\frac{dP_t}{P_t} = \mu dt + \sigma_P dZ_t^P. \tag{7}$$

The dynamics of the ESG fund price S_t follow the stochastic differential equation:

$$\frac{dS_t}{S_t} = (\bar{f} - \theta_S \bar{g}_t)dt + \sigma_S dZ_t^S \tag{8}$$

$$d\langle Z_t^S, Z_t^P \rangle = \rho_M dt \tag{9}$$

The process $\bar{g}_t > 0$ denotes the fund's fundamental sustainability level and is not observable to the investor. This level of sustainability has a direct short-term effect on expected returns $\bar{f} - \theta_S \bar{g}_t$. This effect is scaled by the parameter θ_S. A positive θ_S implies that the greener the fund, the lower its expected returns, in line with Pastor, Stambaugh, and Taylor (2021) and Pedersen, Fitzgibbons, and Pomorski (2021b).[6]

At time 0 and at future random times τ_i, the fund changes its level of sustainability, which is randomly drawn from a discrete set of values $\{\bar{g}_k\}$, $k = 1, ..., 3$. $\bar{g}_1$ (resp. $\bar{g}_3$) represents a low (resp. high) level of sustainability. We remain agnostic as to what causes a change in $\bar{g}_t$. It can either be a change in the holdings of the fund – actively managed funds rebalance their portfolio, potentially changing the overall sustainability level of the fund – or a change in the sustainability level of firms in the fund – firms that decide, for example, to pursue or abandon green investments –. Between the change times τ_i, $\bar{g}_t$ is constant. Change times are generated by a Poisson process.

[5]We consider an information structure that obeys standard assumptions. There exists a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ supporting the vector of four Brownian motions $Z_t = (Z_t^P, Z_t^S, Z_t^g)$, and two Poisson processes $M_t = (M_t^{\bar{g}}, M_t^g)$. $\mathcal{P}$ is the corresponding measure and $\mathcal{F}$ is a right-continuous increasing filtration generated by $Z \times M$. Following Dybvig and Huang (1988) and Cox and Huang (1989), we restrict the set of admissible strategies to those that satisfy the standard integrability conditions. All policies are appropriately adapted to $\mathcal{F}_t$.

[6]As described in Section 2, realized returns on green assets have been shown to be higher at times than realized returns on brown assets, due to demand pressure among others. However, in line with the theory, expected returns on green stocks have been shown to be lower than expected returns on brown stocks, see, e.g., Pastor, Stambaugh, and Taylor (2021).

The transition matrix of $\bar{g}$ is denoted by Λ, where Λ_{kj} denotes the jump intensity from $\bar{g}_k$ to state $\bar{g}_j$.

The investor estimates the sustainability level of the ESG fund from a signal g_t, e.g., an ESG score. This signal is modelled as a mean-reverting process, with long-term mean equal to the fund's sustainability level.

$$dg_t = -\kappa_g(g_t - \bar{g}_t)dt + \sigma_g dZ_t^g + dJ_t^g. \tag{10}$$

The noise of the signal is captured by the volatility parameter σ_g. This noise is due to the challenges one faces when measuring the sustainability level of a firm. We refer to the uncertainty of the signal coming from this term as *measurement uncertainty*. The mean reversion speed parameter κ_g allows the signal to be temporarily biased, e.g., as a result of the firm's disclosure policy. Greenwashing would, for example, pull the ESG score above its true level. Lack of disclosure would pull it below. We refer to the uncertainty caused by this term as *greenwashing uncertainty*.

The jump term dJ_t^g allows the signal to suddenly revert to the true sustainability level g_t. Such shift occurs, for example, when a firm is caught greenwashing. Jump times are driven by a Poisson process, whose intensity is dependent on the gap between the signal and the true sustainability level, $g - \bar{g}$. The larger this gap, the higher the probability of an incident. Upon a jump, the signal g resets at $\bar{g}_t$. That is, measurement uncertainty remains, but greenwashing uncertainty is shut down.

The environmental score is not correlated with the expected returns of stocks but is positively correlated with the stock returns:

$$d\langle Z^S, Z^g\rangle_t = \rho_g dt > 0 \tag{11}$$

This feedback effect between shocks in the market environmental score and shocks in stock returns captures the increase in demand for ESG funds that is triggered when ratings increase.

The investor estimates $\bar{g}$ from g_t, as the value g_k on the grid of possible values which is the closest to g_t. We assume that there is no learning.

4.2 Investor's Problem

The investor continuously rebalances their wealth W_t between the three assets, and consumes at rate $c_t = C_t/W_t$. We denote by π_t^S the fraction of wealth allocated to the ESG fund and

π_t^P to the market index.

The investor is of one of three different types, in line with the types uncovered in Section 3. Investor 1 is a *CRRA investor*, with no non-pecuniary preferences:

$$\text{1. CRRA investor:} \quad U^{C\hat{R}RA}(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{12}$$

The risk aversion coefficient γ is such that $\gamma > 1$.

Investor 2 is a *green investor*, with standard non-pecuniary preferences:

$$\text{2. Green investor:} \quad U^{green}(C_t, \pi_t, \hat{g}_t) = \frac{C_t^{1-\gamma}}{1-\gamma}(1 + b_1\, \pi_t^S\, \hat{g}_t)^{1-\gamma}, \tag{13}$$

Investor 2 does not observe the fundamental sustainability level of the ESG funds, and therefore estimates it. Their utility increases monotonically in this estimated level $\hat{g}_t$: $b_1 > 0$. As the fund's estimated sustainability level increases, the second component of the investor's utility, $1 + b_1\, \pi_t\, \hat{g}_t$, becomes smaller and as a result the utility – which takes negative values – increases. In the particular case where the investor observes the firm's ESG level $\bar{g}$, we have $\hat{g} = \bar{g}$, and the investor's value function increases with the true ESG level of the firm. This specification is close to the one used by Avramov et al. (2022), in which $\hat{g} = g$.

Investor 3 is a *threshold investor*, and receives additional utility from having a portfolio whose ESG rating exceeds a given threshold:

$$\text{3. Threshold investor:} \quad U^{threshold}(C_t, \pi_t, g_t) = \frac{C_t^{1-\gamma}}{1-\gamma}(1 + b_2\, \pi_t^S 1_{g_t > \chi})^{1-\gamma}. \tag{14}$$

Whereas the green investor cares about the estimated fundamental sustainability level of the ESG fund, the threshold investor only cares about whether the signal is above a given threshold. This specification is in line with investors paying attention to a label, given to a fraction of companies based on their ESG rating. To keep specifications (13) and (14) comparable, we choose $b_2 = b_1 \bar{g}_2$.

For all types of investors, the evolution of the investor's wealth is given by:

$$\frac{dW_t}{W_t} = \left(r + (\bar{f} - \theta_S \bar{g}_t - r)\pi_t^S + (\mu - r)\pi_t^P - c_t\right) dt + \pi_t^S \sigma_S dZ_t^S + \pi_t^P \sigma_P dZ_t^P. \tag{15}$$

We assume no short-selling, i.e., $\pi^S, \pi^P \geq 0$.

We assume that investors have infinite horizon, and that the value function is given by

$$J(W_t, f_t, g_t, \bar{g}_t) = \max_{\{\pi^S, \pi^P, c\}} \mathbb{E}_t \left[\int_t^\infty e^{-\beta(u-t)} U(C_u, .) du \right], \tag{16}$$

where $U(C_u, .)$ is one of the three utility functions defined in (12), (13) and (14), subject to (15). The parameter β denotes the investor's subjective discount factor.

The investor's value function, optimal consumption and allocation solve the Hamilton-Jacobi-Bellman (HJB) equation given in Appendix A. The value function can be written as the product of a power function of wealth and a function of the wealth composition:[7]

$$F(W) = W^{1-\gamma} H(\pi). \tag{17}$$

The function H is characterized by the set of ODEs shown in Appendix A. Our method for generating numerical results is also detailed in Appendix A.

4.3 Discussion on Model Assumptions

Our model involves a number of assumptions, which we discuss below.

First, we assume that the fund's sustainability level has an effect on its expected returns, but not on its volatility. Although green assets are commonly believed to have lower risk, Na (2024) shows that this fact only holds for firms which aim to increase earnings while becoming greener. Instead, he finds that firms which focus on sustainability without trying to "grow the pie" (Edmans (2022)) increase their risk.

First, we model the times at which a fund's sustainability level chances with a Poisson process. In practice, companies change their greenness by investing in climate-related actions, e.g., lowering their carbon emissions or better dealing with their waste. These decisions are taken at times that are most likely not random. Similarly, asset managers rebalance funds' holdings based on market conditions and on expectations. However, neither the motives nor the timing of the changes are known by investors in the fund. Our specification is similar to the government policy specification of Pastor and Veronesi (2012), who model the time of policy change as exogenous. In this paper, we take the perspective of the investor, who does not know when changes in ESG funds' sustainability will occur and sees these times as random.

[7] Because the utility function is homothetic and the return processes have constant moments, the value function F is homogenous of degree $1-\gamma$ in wealth W.

Second, we assume that the investor does not attempt to learn the true sustainability level of the ESG fund from the observed signal. Developing in-house tools to evaluate firms' sustainability levels is not only costly for a firm but also challenging. Standard ratings use public information, and for some of them additional information they receive from surveys sent to firms, or interviews with management. Most institutional investors do not have the resources to collect more information than is already used by rating agencies. Even for the most sophisticated ones, collecting additional information would not be sufficient, as this information would need to be checked and updated regularly.

5 Model Calibration

In this section, we detail our model calibration. Without loss of generality, we assume that the fundamental sustainability level of the ESG fund, $\bar{g}_t$, can take the values $\bar{g}_1 = 0.25$, $\bar{g}_2 = 0.75$ and $\bar{g}_3 = 1.25$.

5.1 Market fund dynamics

Unlike for ESG funds, public equities enjoy a long history of data. We use the past thirty years of data to calibrate the market dynamics. The average 3-month Treasury bill is $r = 0.03$. The mean and volatility of the S&P 500 index log returns are 0.08, and 0.15. Therefore we use $\mu = \bar{f} = 0.08$ and $\sigma_P = 0.15$.

5.2 ESG fund dynamics

We calibrate the parameters of the dynamics of the ESG fund's return by matching the distribution of returns as implied by the model to the empirical distribution of green fund returns.

We use Morningstar fund data. The Morningstar Sustainability Rating, or globe rating, is a monthly measure of how sustainable a fund's holdings are relative to peers. An aggregated score is calculated for each fund by taking a weighted average of the ESG ratings of their holdings, obtained from Sustainalytics. Funds are then ranked on the basis of these scores. They are given five globes (four, three, two and one respectively) if they are in the top 10% of funds within their category (between 10% and 32.5%, between 32.5% and 67.5%, between 67.5% and 90%, and above 90%).

We extract the monthly returns and Globes ratings of all funds in the Morningstar

database with a lifespan of more than three years. We drop all funds that are not active at the end of the time period of the Morningstar dataset and funds that have missing data in their return. We then calculate the average number of globes of each fund over the latest two years of data, from June 2020 to June 2022. In total, the data set contains 5143 funds, 540 of which have a time series average rating larger than or equal to 4. We then sort all the funds by their average rating, consider those with an average rating of 4 or more, and divide them into three groups of equal sizes. As all of these funds have a sustainability rating that is above average, we consider them all to qualify as candidates for ESG funds. We use these three groups to estimate the parameters driving the dynamics of the ESG fund for the three different values $\bar{g}_k$. For each group of fund, we compute the empirical cumulative distribution function of fund returns over two years. The average two-year fund returns in the greenest group is 13.6%, it is 16% in the second greenest and 16.4% in the third group.[8]

The model-implied cumulative distribution function of fund returns is obtained using Monte-Carlo simulations of the discretized equation (8). We estimate the values of $\bar{f}, \sigma_S$ and θ_S by minimizing the average least-squares distance between the model-implied and the empirical cumulative distribution functions. We give equal weight to the three least-square distances as the three groups of funds contain the same number of funds. As shown in Figure 3, the model matches the data remarkably well. We obtain $\bar{f} = 0.08, \sigma_S = 0.15$, and $\theta_S = 0.007$. Interestingly the first two moments of the ESG fund, setting $\bar{g}_t = 0$, are equal to the moments of the market fund. The estimated value of θ_S implies that over a year, on average, the greenest funds underperform the market by 0.9%. This value is in line with the interval found by Pastor et al. (2022), between 0.4% and 2.4%.

The correlation between market fund returns and ESG fund returns is set at $\rho_M = 0.95$, in line with ESG funds tracking the same portfolio as the main index, under additional sustainability constraints.

[8]Until 2021, green funds exhibited higher realized returns than brown funds. Pastor et al. (2022) and van der Beck (2022) show that these higher returns are explained by increasing climate concerns leading to demand pressure for green assets. After isolating the demand pressure component, they find that green assets underperformed brown assets. The climate news index of Gourier and Mathurin (2024) shows that climate concerns have decreased since the end of 2021, which is consistent with the sharp decrease in sustainable fund flows reported by Morningstar from Q1 2022. See, e.g., https://www.morningstar.com/sustainable-investing/us-esg-funds-bleed-less-money-q2-2024. To reduce the effect of demand pressure on our estimate of expected returns, we focus on the last two years of available data in our estimation of ESG fund parameters.

5.3 Dynamics of the signal

We calibrate the dynamics of the ESG signal by matching the average cumulative distribution function of funds' Globe ratings over a three-year period, from June 2019 to June 2022, to the model-implied average cumulative distribution function without the jump term J_t.[9] The average is taken over the three groups of funds defined in Section 5.2. For each group, the signal follows an Ornstein-Uhlenbeck process. We use as initial value the average rating of funds in each group in June 2019. We minimize the Euclidean distance between the data implied and model implied cumulative distribution functions. The fit of the model is shown in Figure 4. We obtain $\kappa_g = 0.04$ and $\sigma_g = 0.12$. The low value of κ_g, which implies a half-life of more than 17 years, shows that ratings are highly persistent.

We parameterize the Poisson process driving the jump term J_t such that a jump is more likely to occur when the gap between the signal and the true sustainability of the ESG fund is large. Specifically, the Poisson process has intensity $\nu = \nu_1$ whenever the estimated sustainability $\hat{g}_t$ is one level on the grid higher than the true sustainability level $\bar{g}_t$, and $\nu = \nu_2$ whenever it is two levels higher:

$$\nu = \nu_1 \cdot 1_{0.25<g_t-\bar{g}_t\leq 0.75} + \nu_2 \cdot 1_{g_t-\bar{g}_t>0.75} \tag{18}$$

We estimate ν_1 and ν_2 using environmental incidents from RepRisk. Statistics on reported incidents among firms are reported in Table 3. The percentage of firms with reported incidents increased over the years. We filter out non-US-based firms and firms that do not have an ESG rating in 2021. We compute the median ESG rating of all firms, and consider that all firms with an ESG rating above the median have a green signal. We estimate ν_1 as the ratio of firms with a green signal that had a reported incident in 2021. Among these firms, those with an ESG rating in the highest tercile are used to estimate ν_2. We obtain $\nu_1 = 0.3$ and $\nu_2 = 0.4$. Firms whose ratings deviate from their true sustainability level by one rating therefore have a probability of 30% to be reported, and firms whose ratings deviate by two levels have a probability of 70% to be exposed within a year.

To estimate the correlation between the signal and ESG fund returns, we calculate the average correlation between the Globes ratings and returns of funds with a rating of 4 Globes or more. We obtain $\rho_g = 0.02$. On average, Globe ratings and returns on green funds are therefore nearly uncorrelated.

[9] Incident reports are rare events, therefore we calibrate them separately.

5.4 Investors' preferences and discounting

We use standard values for the risk aversion and discount factor: $\gamma = 4$ and $\beta = 0.1$.

The parameter driving non-pecuniary utility, b_1, needs to be large enough so that non-pecuniary utility outweighs the cost of the fund when the level of sustainability is high. We choose accordingly $b_1 = 0.2$.

All our parameters are listed in Table 4.

6 Empirical Results

In this part, we present the optimal portfolio allocations and consumption policies of the three types of investors: CRRA investors (small and medium-sized institutional investors), green investors (large insurers) and threshold investors (large banks, investment companies and advisors). We compare, for each type of investor, these optimal allocations and policies to those of economies without measurement and greenwashing uncertainty.

6.1 Optimal Portfolio and Consumption

6.1.1 CRRA investors

We first consider investors with CRRA utility given in (12).

Panel A of Figure 5 displays the investor's ESG fund and market index allocation as functions of the ESG signal. If the ESG signal is low (between 0 and 0.5), the investor estimates the firm's sustainability level at 0.25, and invests 15% of wealth in the ESG fund, 45% in the market index, and 40% in the bond. As the expected returns of the ESG fund are lower than those of the market index, the positive allocation to the ESG fund comes from the imperfect correlation between the two assets. The investor faces a trade-off between having a more diversified portfolio, and paying the cost of being greener. When the signal is above 0.5, the investor estimates the true sustainability level at 0.75 or at 1.25, and does no longer allocate any wealth to the ESG fund, as the cost of being green dominates the benefits of diversification. The allocation to the market index increases to 55%, and the allocation to bonds to 45%.

Panel B shows that the optimal allocation corresponds to an investor's green tilt at 2% when the signal is low, and at 0 when for larger values of the signal. This tilt is consistent in shape with the green tilt of small and medium-sized investors described in Figure 2, Panel

B.

When the ESG signal is close to the fundamental sustainability level of the fund, there is no misallocation. Misallocation only occurs when the signal is far from the true level, and investors misestimate the moments of returns: if the signal is too high, the investor under-allocates to the ESG fund, and if it is too low, they over-allocate to it. This misallocation translates into changes in the optimal consumption policy: if the signal is too high (respectively, too low), the investor consumes less (respectively, more) than the optimal amount. However, the scale of Panel C shows that this errors in consumption are small.

It is therefore not surprising that the loss of welfare, due to the effect of not observing $\bar{g}$, is small as well irrespective of the fundamental fund's sustainability level $\bar{g}_t$, and of the signal value g_t. Panels D to F compare the investor's welfare when observing (solid line) and when not observing (dashed line) $\bar{g}$, for a true sustainability level of $\bar{g}_t = \bar{g}_1$ (Panel D), $\bar{g}_t = \bar{g}_2$ (Panel E) and $\bar{g}_t = \bar{g}_3$ (Panel F).

6.1.2 Green investors

We now turn to green investors, who have a non-pecuniary taste for high ESG standards according to (13). The results are given in Figure 7.

As Panel A shows, the green investor optimally allocates much more to the ESG fund than the CRRA investor: between 40 and 65% of wealth against no more than 15%. Furthermore, the allocation increases with the ESG signal. The utility drawn from having a green portfolio outweighs the utility lost from the lower return of the ESG fund. Accordingly, the allocation to the market index is small, less than 20% of wealth, and decreases with the ESG signal. Panel B shows that the increase in ESG fund allocation translates into a investor's green tilt that is much higher than the one of the CRRA investor, increasing from 8% to 13% depending on the ESG rating. This pattern is in line with the green tilt of insurers observed in Figure 2, Panel A.

Panel C reveals that not observing the firm's ESG level has a larger effect on the green investor's consumption whenever the signal is far from the true sustainability level. Mirroring the change in market index allocation, an investor who estimates a higher ESG level for the fund decreases their consumption. The non-pecuniary utility substitutes for the utility from consumption, and the substitution effect reduces the optimal consumption rate. If the ESG signal is high but the fund's true sustainability level is low, the investor under-consumes. Similarly, they over-consume if the ESG signal is low but the true sustainability level is high.

Panel D reveals that welfare increases with the ESG signal. For low values of the ESG

signal, welfare is much lower than if the investor were to observe the true sustainability level, but for high values of the signal, the two levels of welfare are close.

6.1.3 Threshold investors

Figure 9 illustrates the optimal portfolio and consumption policy of a threshold investor, with utility function given by (14).

When the ESG rating is low, the investor's demand for the ESG fund is the same as the one of a CRRA investor, about one third the one of a green investor (Panel A). This is because when the estimated sustainability level is below the threshold, the investor behaves as a CRRA investor. When the estimated ESG level crosses the threshold, the investor gains additional non-pecuniary utility and increases the allocation to the ESG fund to 60%, similar to the green investor. For larger values of the ESG signal, the threshold investor does not gain additional utility, and bears the estimated cost of greenness. The allocation to the ESG fund thus drops, from 60% to 35% as the ESG signal exceeds 1. The corresponding green tilt of the threshold investor is around 2% for $g < 0.5$, 12% for g between 0.5 and 1, and 7% for g larger than 1. This tilt is close to the tilt of 6% reported by Pastor et al. (2023).

Furthermore, as shown on Panel C, the ESG signal only has a large impact on the optimal consumption rate if the signal suggests a level of sustainability on one side of the threshold, whereas the true level is on the other side. The investor consumes more when the ESG signal is below the threshold. Above the threshold, the allocation to the ESG fund increases and as a result, the investor consumes less.

Panels D to F compare the welfare of a threshold investor to the case in which they observe the true sustainability level. As for the CRRA and the green investors, the loss of welfare is the largest for small values of the ESG signal. This loss of welfare is however smaller than for the threshold investor. As the ESG signal increases, it becomes smaller.

6.2 Cost of ESG Rating Uncertainty

To compare economies, we follow Gourier, Phalippou, and Westerfield (2024) and define the welfare cost of economy A relative to economy B to be the fraction of wealth ϕ that investors are willing to give up to switch from A to B. For example, the welfare cost of our main economy, with measurement and greenwashing uncertainty, relative to an economy with no measurement and greenwashing uncertainty, is the fraction of wealth investors are willing to pay to observe the value of the fund's sustainability level $\bar{g}_t$ at all times. Formally, ϕ is given

by

$$(1-\phi)^{1-\gamma}H^B = H^A \tag{19}$$

where H^A, H^B represent the value functions of economy A and B respectively. The welfare cost varies between types of investors.

We also compute the return premium of economy A with respect to economy B, as the additional return of the ESG fund that would be needed in economy A to make the investor indifferent between the two economies.[10]

We define two economies. **Economy 0** is our baseline economy, in which there is ESG rating uncertainty. **Economy 1** sets the volatility of the signal to zero: $\sigma_g = 0$, and is hence an economy with no measurement uncertainty. **Economy 2** sets $g = \bar{g}$, hence the signal is observed and there is neither measurement nor greenwashing uncertainty.

6.2.1 CRRA investors

Figure 6 displays the welfare cost and the return premium associated with measurement uncertainty, and with total ESG rating uncertainty (measurement and greenwashing). The cost of measurement uncertainty is small, irrespective of the true sustainability level $\bar{g}$. The maximum welfare cost is reached for a signal of $g = 0.5$. This is indeed the threshold at which the difference of allocation below and above is the highest. Volatility makes investors switch between these two allocations, with the risk of being on the wrong side of 0.5. But even at this threshold, the welfare cost (Panel A) is at 0.1% of wealth, and the return premium (Panel B) at 0.015%.

The cost of greenwashing uncertainty is larger, especially for small values of the signal. As a result, the total cost of ESG uncertainty is monotonically decreasing, reaching 0.45% of welfare cost (Panel C) and 0.07% of return premium (Panel D). At low values of the signal, the investor estimates that the fund's sustainability level is the lowest value of the grid, i.e., 0.25. With greenwashing, there is a chance that the true sustainability level is in fact larger, with higher cost and lower expected return than what the investor anticipates.

Overall, these results suggest that the uncertainty of ESG ratings does not matter much for CRRA investors, i.e., for small and medium-sized institutional investors.

[10] As noted by Gourier, Phalippou, and Westerfield (2024), the two welfare measures should be interpreted differently. The welfare cost is a one-time payment to switch economies, so it is strongly increasing in the optimal ESG fund allocation. In contrast, the return premium impacts the investor in proportion to the amount allocated to the ESG fund, so it is closer to a *per unit* cost.

6.2.2 Green investors

Figure 8 shows that the welfare cost associated with measurement uncertainty (Panel A) is relatively small but still twice the one for a CRRA investor, at 0.8% of wealth for $g = 0.5$ and $g = 1$. As for the CRRA investor, these values are threshold that are important for the estimation of $\bar{g}$: below 0.5 (respectively, 1), investors estimate $\hat{g} = 0.25$ (respectively, 0.75), and above, they estimate $\hat{g} = 0.75$ and 1.25. These estimates change both the portfolio weights and the consumption level. The return premium (Panel B) reaches a maximum of 0.6% which is five times larger than the maximum of the CRRA investor.

Greenwashing uncertainty also has a large welfare cost (Panel C): 3% to 4% of wealth, for low values of the ESG signal. Similarly, the return premium (Panel D) is nearly 2%. For large values of the ESG signal, both welfare cost and return premium for greenwashing uncertainty are around 1%.

Both measurement and greenwashing uncertainty thus matter to green investors (insurers), who have incentives to pay to remove these frictions from the economy.

6.2.3 Threshold investors

Figure 10, Panels A and B, show that the threshold investor has a welfare cost that reaches about the same maximum as a green investor, but a return premium that is less than half. Similarly, the welfare cost of greenwashing uncertainty (Panel C) is about half the one of a green investor, at nearly 2% for low values of the ESG signal. The corresponding return premium (Panel D) is four times smaller, due to the smaller allocation to the ESG fund.

Threshold investors, who represent the vast majority of investors, therefore have much less incentives than green investors to push towards more transparent ESG signals.

7 Robustness tests

7.1 Model extension

We consider an extension of the model in which the expected returns are driven by the mean-reverting process f_t, whose dynamics follow

$$df_t = \kappa_f[(\bar{f} + \theta_f \bar{g}_t) - f_t]dt + \sigma_f dZ_t^f \tag{20}$$

The Brownian motions Z_t^P and Z_t^f are assumed to be independent.

This extension has a insignificant effect on our results.

References

Aboody, D. and B. Lev (1998). The value relevance of intangibles: the case of software capitalization. *Journal of Accounting Research 36*, 161–191.

Albuquerque, R., Y. Koskinen, and C. Zhang (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management Science 65*(10), 4451–4469.

Avramov, D., S. Cheng, A. Lioui, and A. Tarelli (2022). Sustainable investing with esg rating uncertainty. *Journal of Financial Economics 145*, 642–664.

Baldauf, M., L. Garlappi, and C. Yannelis (2020). Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies 33*(3), 1256–1295.

Berg, F., J. F. Koelbel, A. Pavlova, and R. Rigobon (2023). ESG Confusion and Stock Returns: Tackling the Problem of Noise. Working Paper.

Berg, F., J. F. Koelbel, and R. Rigobon (2022). Aggregate Confusion: The Divergence of ESG Ratings. *Review of Finance 26*(6), 1315–1344.

Berg, F., J. F. Kölbel, A. Pavlova, and R. Rigobon (2024). Esg confusion and stock returns: Tackling the problem of noise. Working Paper.

Bernstein, A., M. T. Gustafson, and R. Lewis (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics 134*(2), 253–272.

BlackRock (2020). Global Sustainable Investing Survey. Technical report.

Bolton, P. and M. T. Kacperczyk (2021). Do investors care about carbon risk? *Journal of Financial Economics 142*, 517–549.

Bolton, P. and M. T. Kacperczyk (2023). Global pricing of carbon-transition risk. *Journal of Finance 78*, 3677–3754.

Cartellier, F., P. Tankov, and D. Zerbib (2024). Can investors curb greenwashing? Working Paper.

Chan, L., J. Lakonishok, and T. Sougiannis (2001). The stock market valuation of research and development expenditures. *Journal of Finance 56*, 2431–2456.

Chang, G. (2014). Alternative formulation of the kalman filter for correlated process and observation noise. *IET science, measurement & technology 8*(5), 310–318.

Chava, S. (2014). Environmental externalities and cost of capital. *Management Science 60*(9), 2223–2247.

Christensen, D. M., G. Serafeim, and A. Sikochi (2022). Why is corporate virtue in the eye of the beholder? the case of esg ratings. *The Accounting Review 97*(1), 147–175.

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

De Angelis, T., P. Tankov, and O. D. Zerbib (2023). Climate impact investing. *Management Science 69*(12), 7669–7692.

Deng, Z., B. Lev, and F. Narin (1999). Science and technology as predictors of stock performance. *Financial Analysts Journal 55*, 20–32.

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

Edmans, A. (2011). Does the stock market fully value intangibles? Employee satisfaction and equity prices. *Journal of Financial Economics 3*(101), 621–640.

Edmans, A. (2022). *Grow the Pie: How Great Companies Deliver Both Purpose and Profit*.

Faccini, R., R. Matin, and G. S. Skiadopoulos (2023). Dissecting climate risks: Are they reflected in stock prices? *Journal of Banking and Finance*. Forthcoming.

Fama, E. F. and K. R. French (2007). Disagreement, tastes, and asset prices. *Journal of financial economics 83*(3), 667–689.

Friedman, H. L. and M. S. Heinle (2016). Taste, information, and asset prices: implications for the valuation of csr. *Review of Accounting Studies 21*, 740–767.

Glossner, S. (2021). Repeat offenders: Esg incident recidivism and investor underreaction. Working Paper.

Goldsmith-Pinkham, P., M. T. Gustafson, R. C. Lewis, and M. Schwert (2023, 05). Sea-Level Rise Exposure and Municipal Bond Yields. *The Review of Financial Studies*.

Görgen, M., A. Jacob, M. Nerlinger, R. Riordan, M. Rohleder, and M. Wilkens (2020). Carbon risk. Working Paper.

Gourier, E. and H. Mathurin (2024). A Greenwashing Index. Working paper.

Gourier, E., L. Phalippou, and M. Westerfield (2024). Capital commitment. *Journal of Finance Forthcoming*.

Heinkel, R., A. Kraus, and J. Zechner (2001). The Effect of Green Investment on Corporate Behavior. *Journal of Financial and Quantitative Analysis 36*(4), 431–449.

Holmes, E. (2003). Kalman filtering for maximum likelihood estimation given corrupted observations. *Natl. Mar. Fish. Serv 22*, 929–932.

Hong, H. and M. Kacperczyk (2009). The price of sin: the effects of social norms on markets. *Journal of Financial Economics 2*(93), 15–36.

Hong, H., F. W. Li, and J. Xu (2019). Climate risks and market efficiency. *Journal of Econometrics 208*(1), 265–281.

Hsu, P.-H., K. Li, and C.-Y. Tsou (2022). The pollution premium. *Journal of Finance*. Forthcoming.

Huynh, T. D. and Y. Xia (2021). Climate change news risk and corporate bond returns. *Journal of Financial and Quantitative Analysis 56*(6), 1985–2009.

Ilhan, E., Z. Sautner, and G. Vilkov (2021). Carbon tail risk. *Review of Financial Studies 34*(3), 1540–1571.

Khan, M., G. Serafeim, and A. Yoon (2016). Corporate Sustainability: First Evidence on Materiality. *The Accounting Review 91*(6), 1697–1724.

Krueger, P., R. G. Brandon, and P. Schmidt (2021). Esg rating disagreement and stock returns. *Financial Analysts Journal 77*(4), 104–127.

Landier, A. and S. Lovo (2023). ESG Investing: How to Optimize Impact? HEC Paris Research Paper No. FIN-2020-1363.

Leippold, M. and T. Yu (2023). The green innovation premium: Evidence from us patents and the stock market. *Swiss Finance Institute Research Paper* (21).

Lev, B. and T. Sougiannis (1996). The capitalization, amortization, and value relevance of R&D. *Journal of Accounting and Economics 21*, 107–138.

Lins, K. V., H. Servaes, and A. Tamayo (2017). Social Capital, Trust, and Firm Performance: The Value of Corporate Social Responsibility during the Financial Crisis. *Journal of Finance 72*(4), 1785–1824.

Na, M. (2024). The silver lining of investor disagreement: Reducing risk with.

Oehmke, M. and M. M. Opp (2024). A theory of socially responsible investment. *Review of Economic Studies*.

Painter, M. (2020). An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics 135*(2), 468–482.

Paribas, B. (2019). The ESG Global Survey. Technical report.

Pastor, L., R. F. Stambaugh, and L. A. Taylor (2023). Green tilts. Technical report, National Bureau of Economic Research.

Pastor, L. S., R. F. Stambaugh, and L. A. Taylor (2021). Sustainable investing in equilibrium. *Journal of Financial Economics 142*(2), 550–571.

Pastor, L. S., R. F. Stambaugh, and L. A. Taylor (2022). Dissecting green returns. *Journal of Financial Economics 146*(2), 403–424.

Pastor, L. S. and P. Veronesi (2012). Uncertainty about Government Policy and Stock Prices. *Journal of Finance 67*(4), 1219–1264.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021a). Responsible investing: The esg-efficient frontier. *Journal of Financial Economics 142*(2), 572–597.

Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021b). Responsible investing: The ESG efficient frontier. *Journal of Financial Economics 142*(1), 572–597.

Schroders (2020). Global Investor Study. Technical report.

Seltzer, L. H., L. S. and Q. Zhu (2022). Climate regulatory risk and corporate bonds. Nanyang Business School Research Paper No. 20-05. FRB of New York Staff Report No. 1014.

Times, F. (2017, September). The ethical investment boom. James Kynge, September 7.

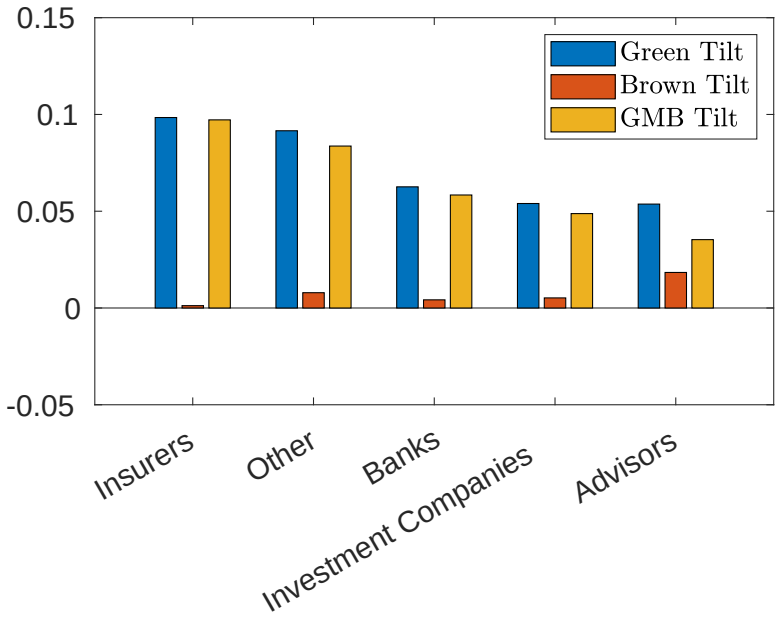
van der Beck, P. (2022). Flow-Driven ESG Returns.

Zerbib, O. D. (2022). A sustainable capital asset pricing model (s-capm): Evidence from environmental integration and sin stock exclusion. *Review of Finance 26*(6), 1345–1388.

Figure 1: Institutional Level Tilts

This figure displays the weighted average green, brown and green-minus-brown (GMB) tilts of groups of investors. The computation of the tilts is described in Section 3.2. In Panel A, investors are grouped by their type as reported by the 13F form: banks, insurers, investment companies, advisors or other. In Panel B, they are split by tercile of their assets under management.

Panel A: Tilts by type of institutional investor



Panel B: Tilts by size

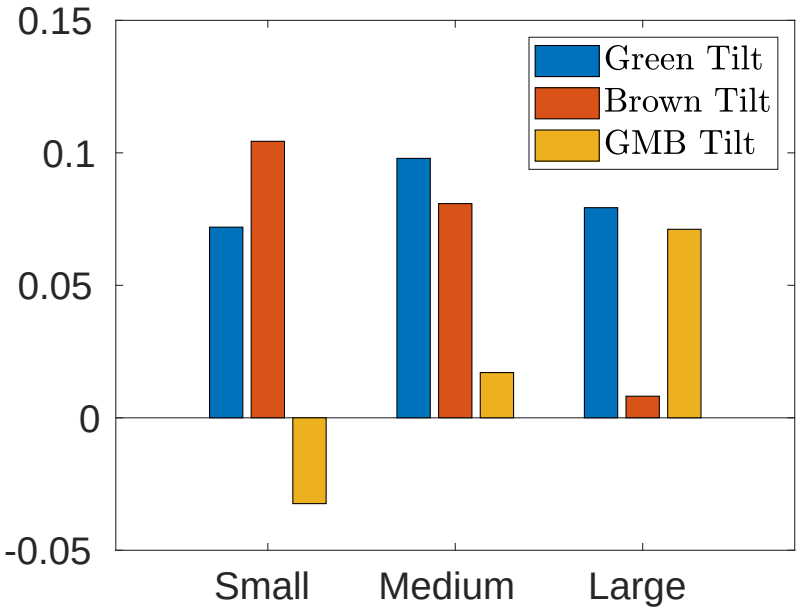
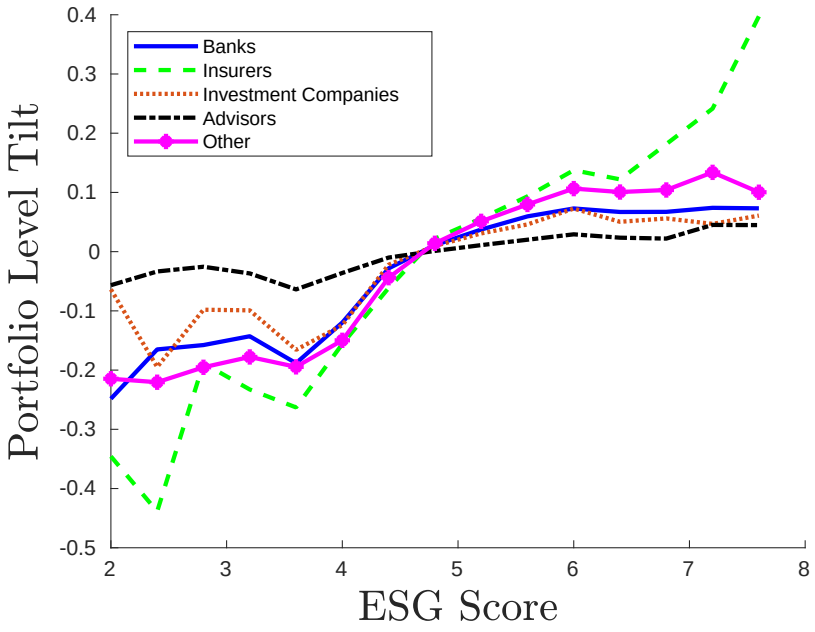


Figure 2: Institutional Level Tilts

This figure displays the weighted average tilts of groups of investors, as functions of assets' ESG scores. The computation of the tilts is described in Section 3.3. In Panel A, investors are grouped by their type as reported by the 13F form: banks, insurers, investment companies, advisors or other. In Panel B, they are split by tercile of their assets under management. X-axis gives the minimum ESG score of each bin.

Panel A: Tilts by type of institutional investor



Panel B: Tilts by size

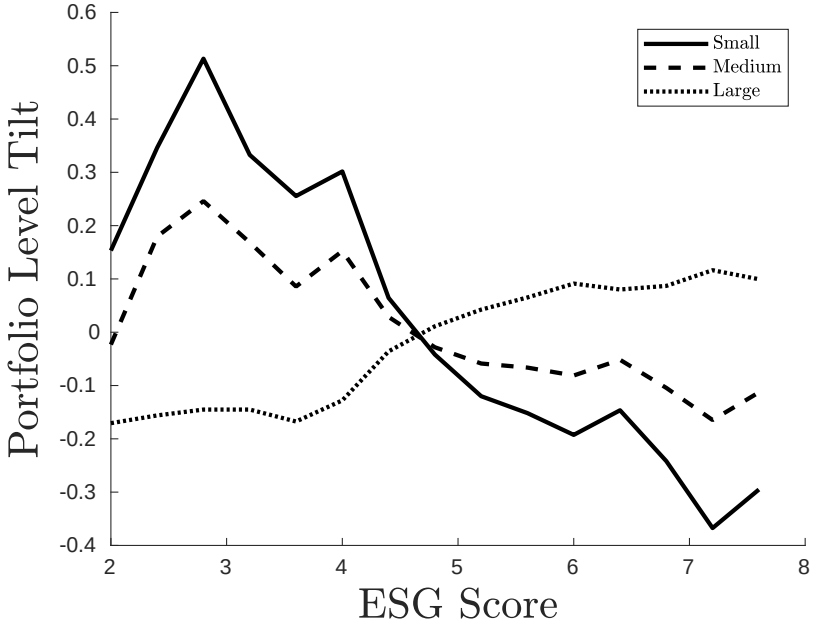


Figure 3: Calibration of ESG Fund Return Dynamics

This figure displays the model implied cumulative distribution function (CDF) of $\ln(S_T/S_0)$ and the CDF implied by data. Here T is equal to two years and S_t is the value of the fund at time t. Panel A, B, C give respectively the fit of the funds in G1, G2 and G3 group. The definition of G1, G2 and G3 group is given in section 5.2. In all panels, the red concrete line gives the model implied CDF and the line of blue crosses gives the data implied CDF.

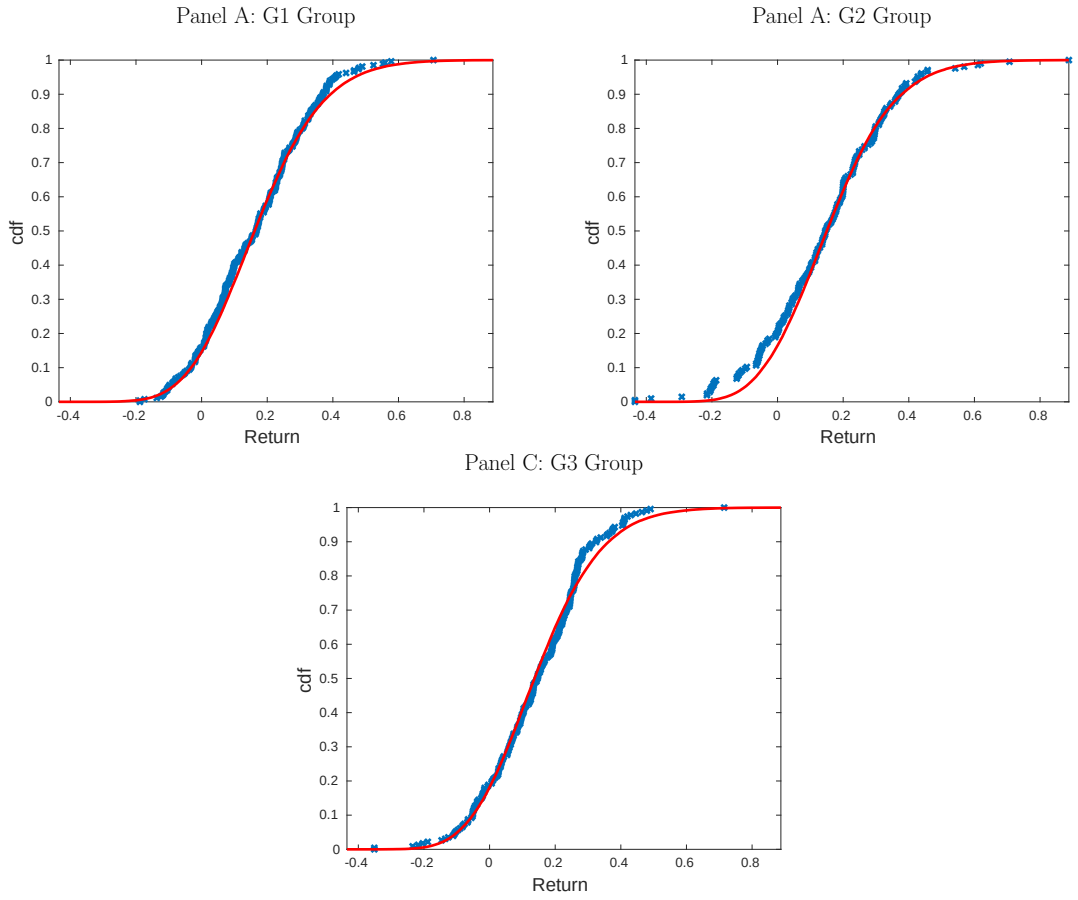


Figure 4: Calibration of ESG Signal

This figure displays the model implied cumulative distribution function (CDF) of the lastest ESG signal g_T and the CDF implied by data. Here T is three years and the signal is represented by the Globes rating of Morningstar. Panel A, B, C give respectively the fit of the funds in G1, G2 and G3 group. The definition of G1, G2 and G3 group is given in section 5.3. In all panels, the orange line gives the model implied CDF and the blue line gives the data implied CDF.

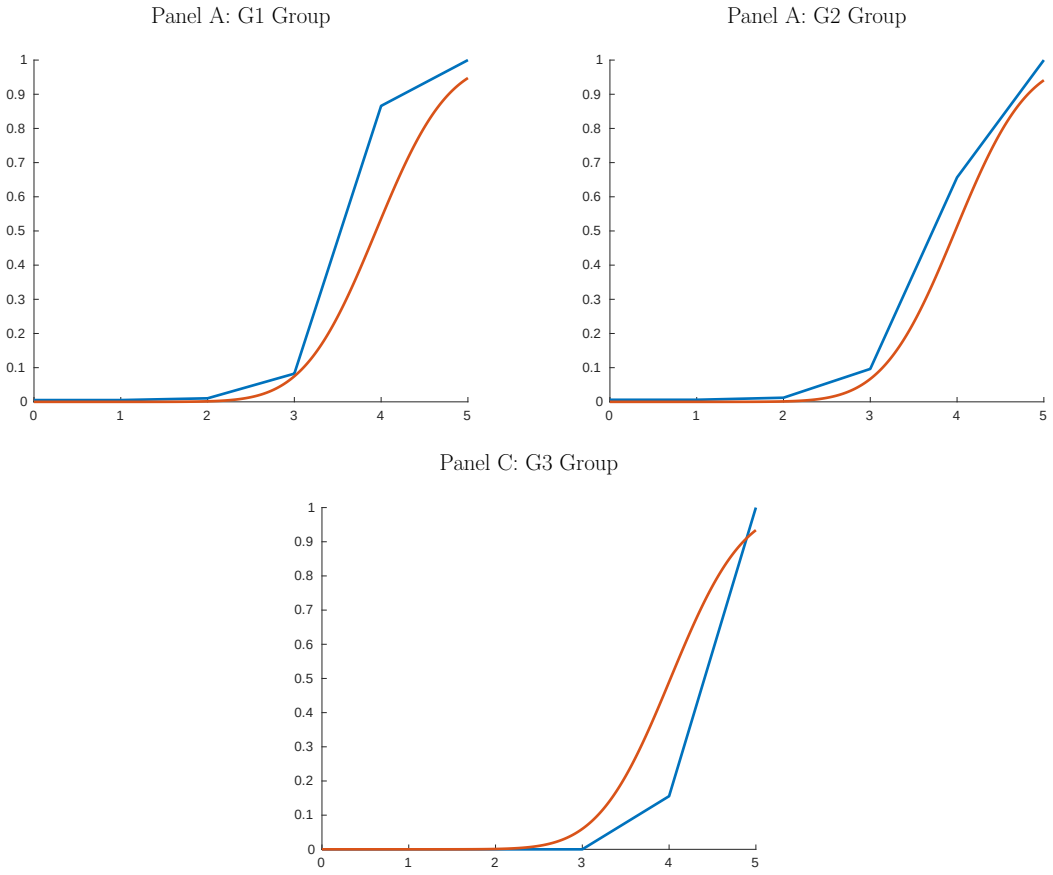


Figure 5: Optimal Allocation and Consumption of the CRRA Investor

Investors have a standard power utility function. They do not observe the true ESG level $\bar{g}$ of the stock. When $g < 0.5$, they infer that $\bar{g} = 0.25$, when $0.5 \leq g < 0.75$ they infer that $\bar{g} = 0.75$, when $0.75 < g$ they infer that $\bar{g} = 1.25$. The optimal portfolio weights and consumption rate are a function of g and does not depend on the value of $\bar{g}$. Panel A illustrates the portfolio weights of the investor. The solid line gives the allocation to the ESG fund and the dashed line gives the allocation to the market index. Panel B shows the green tilt of the investor, as defined in Section 3. Panel C shows the optimal consumption rate. Panel D to F show the welfare of the investor in the case where $\bar{g}$ is equal to 0.25, 0.75 and 1.25 respectively. The solid line represents the case where the investor observed $\bar{g}$ and the dashed line represents the case where the investor does not observe $\bar{g}$.

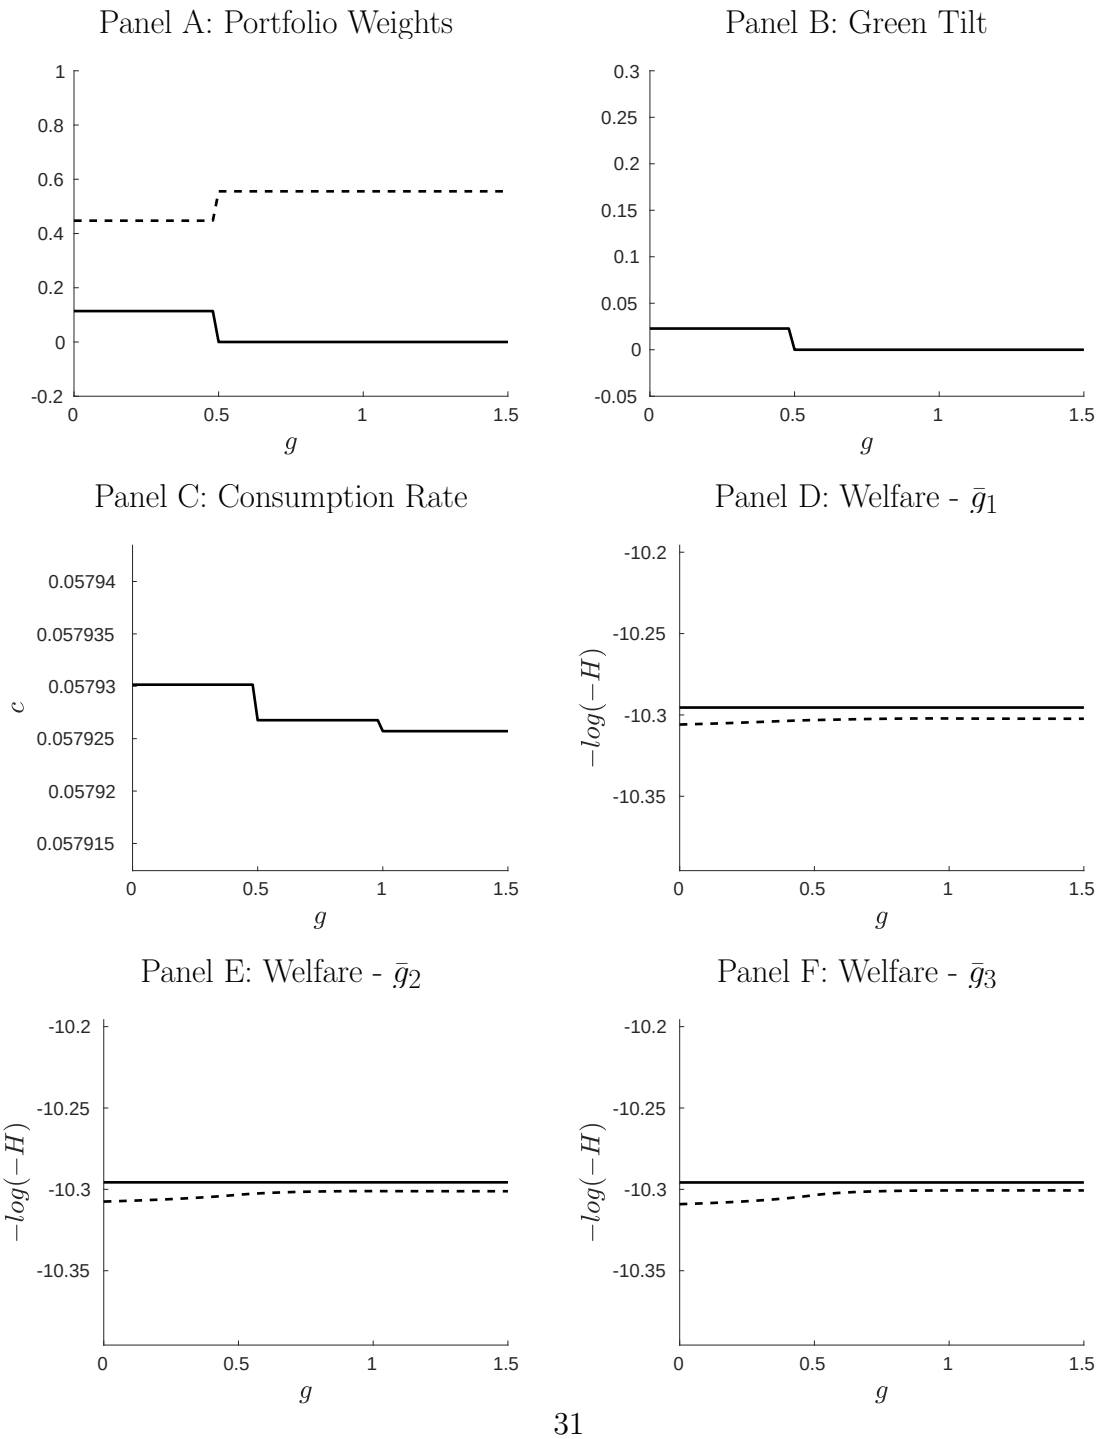
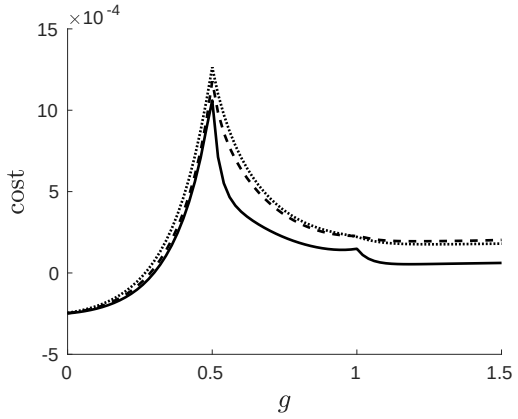


Figure 6: Cost of ESG Fund Level Uncertainty for a CRRA Investor

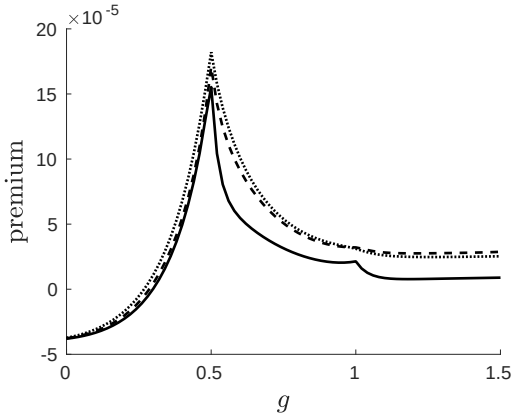
This figure decomposes the cost of ESG rating uncertainty into measurement and greenwashing uncertainty costs. Panels A and B represent, respectively, the welfare cost and the return premium of measurement uncertainty. Economy 1 ($\sigma_g = 0$) is compared to Economy 0 (baseline). Panels C and D display the total welfare cost and return premium of ESG uncertainty (measurement + greenwashing). Economy 2 ($g = \bar{g}$) is compared to Economy 0. The plain, dashed and dotted line correspond to the cases where $\bar{g}$ is equal to $0.25, 0.75$ and 1.25, respectively.

Cost of measurement uncertainty: Economy 1 versus Economy 0

Panel A: Welfare cost

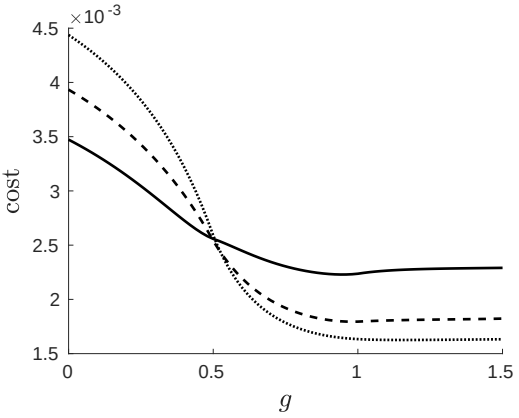


Panel B: Return premium



Cost of ESG rating uncertainty: Economy 2 versus Economy 0

Panel C: Welfare cost



Panel D: Return premium

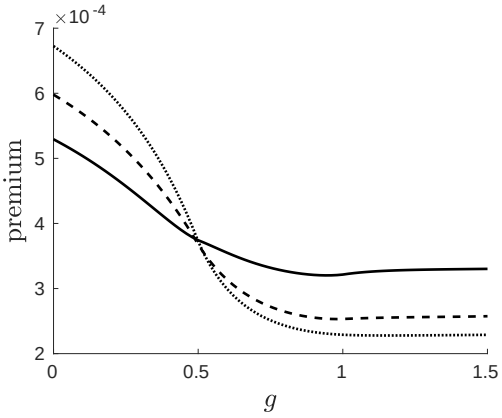


Figure 7: Optimal Allocation and Consumption for a Green Investor

Investors have a taste for high ESG standards ($b = 0.2$, utility function in (13)). They do not observe the true ESG level $\bar{g}$ of the stock. When $g < 0.5$, they infer that $\bar{g} = 0.25$, when $0.5 \leq g < 0.75$ they infer that $\bar{g} = 0.75$, when $0.75 < g$ they infer that $\bar{g} = 1.25$. The optimal portfolio weights and consumption rate are a function of g and does not depend on the value of $\bar{g}$. Panel A illustrates the portfolio weights of the investor. The solid line gives the allocation to the ESG fund and the dashed line gives the allocation to the market index. Panel B shows the green tilt of the investor, as defined in Section 3. Panel C shows the optimal consumption rate. Panel D to F show the welfare of the investor in the case where $\bar{g}$ is equal to 0.25, 0.75 and 1.25 respectively. The solid line represents the case where the investor observed $\bar{g}$ and the dashed line represents the case where the investor does not observe $\bar{g}$.

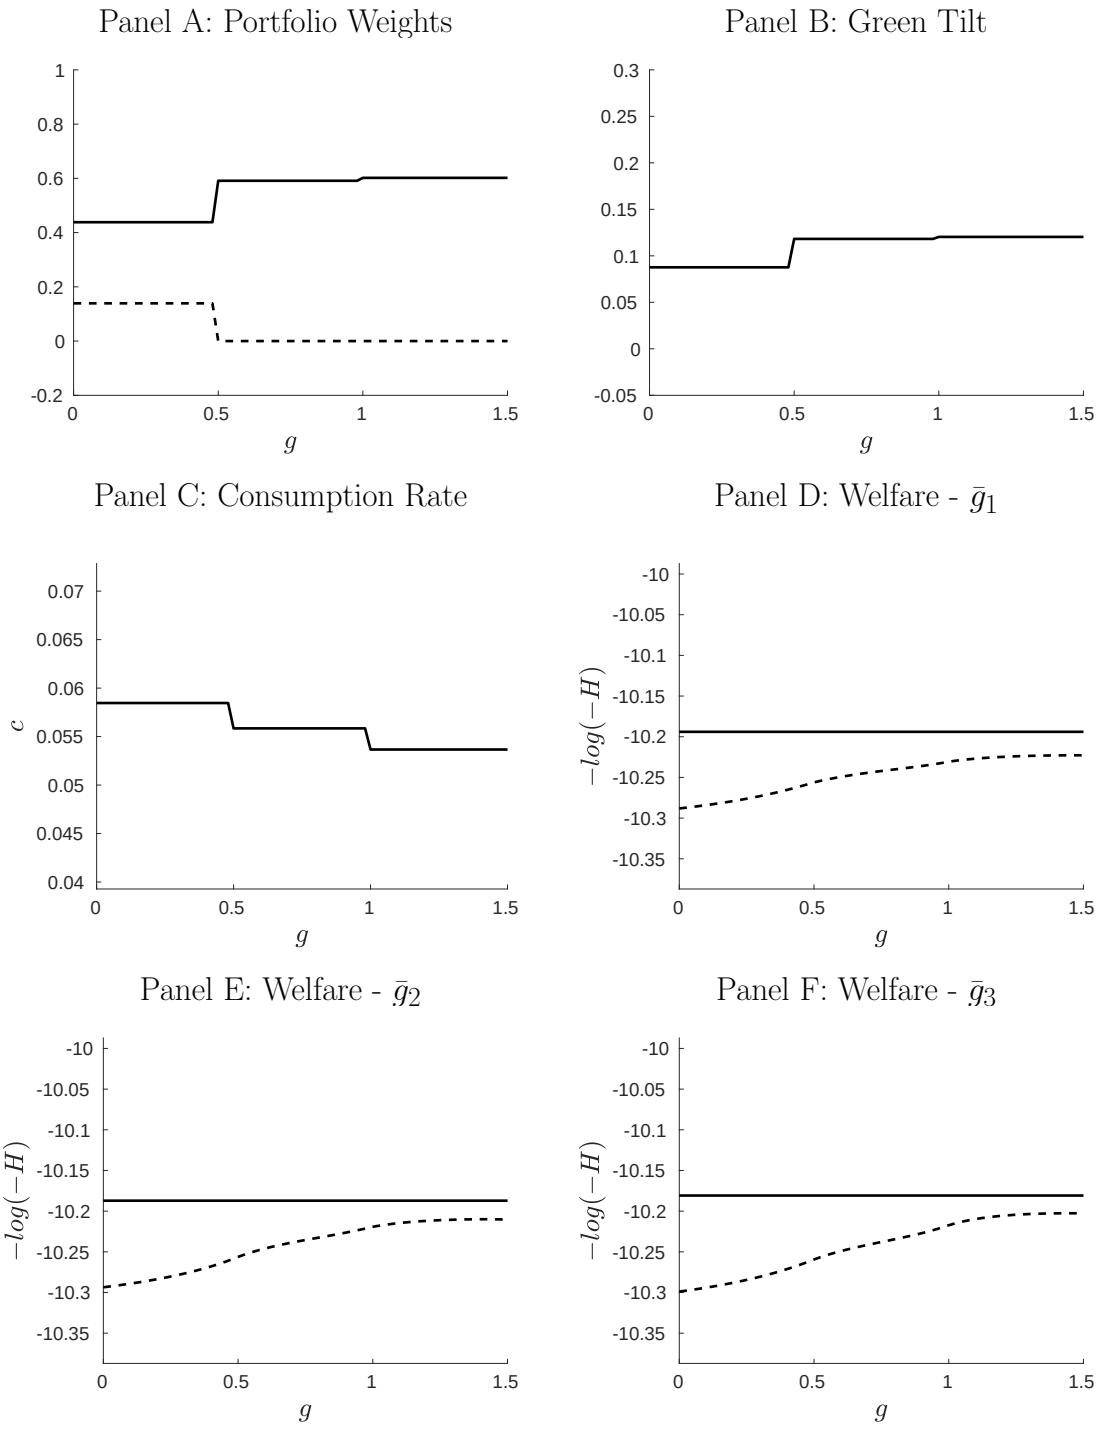
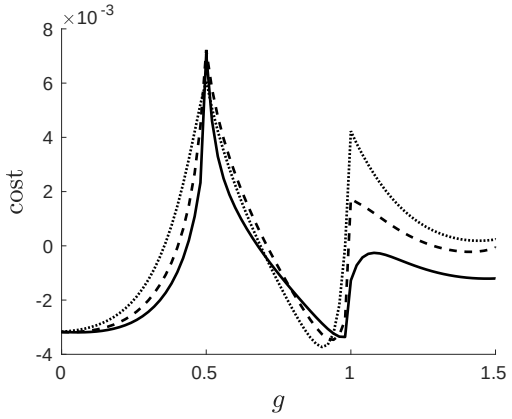


Figure 8: Cost of ESG Fund Level Uncertainty for a Green Investor

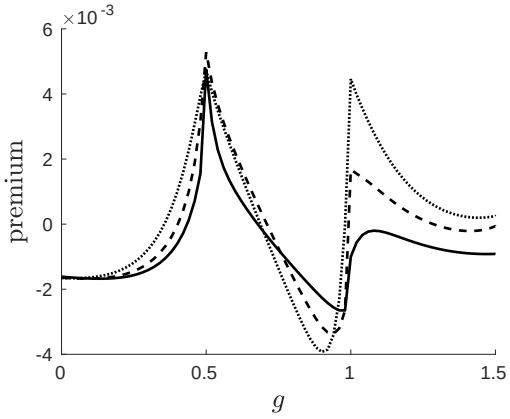
This figure decomposes the cost of ESG rating uncertainty into measurement and greenwashing uncertainty costs, for green investors. Panels A and B represent, respectively, the welfare cost and the return premium of measurement uncertainty. Economy 1 ($\sigma_g = 0$) is compared to Economy 0 (baseline). Panels C and D display the total welfare cost and return premium of ESG uncertainty (measurement + greenwashing). Economy 2 ($g = \bar{g}$) is compared to Economy 0. The plain, dashed and dotted line correspond to the cases where $\bar{g}$ is equal to $0.25, 0.75$ and 1.25, respectively.

Cost of measurement uncertainty: Economy 1 versus Economy 0

Panel A: Welfare cost

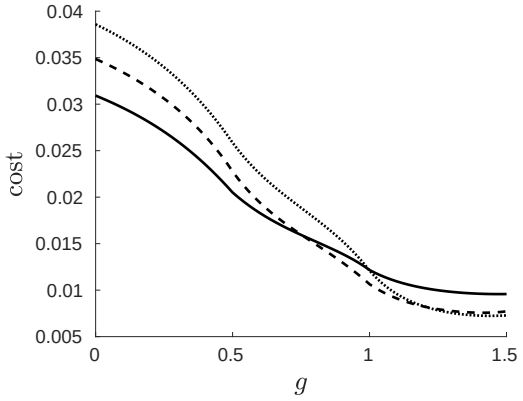


Panel B: Return premium



Cost of ESG rating uncertainty: Economy 2 versus Economy 0

Panel C: Welfare cost



Panel D: Return premium

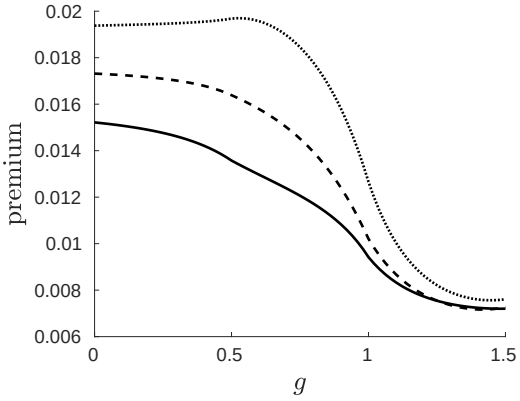


Figure 9: Optimal Allocation and Consumption for Threshold Investor

Investor has a taste for high ESG standards but only cares about the label of the stock ($b = 0.2$, utility function in (14)). The investor does not observe the true stock's ESG level $\bar{g}$. They derive utility from greenness when $g > 0.5$ but their utility from greenness does not increase as g goes up further. The optimal portfolio weights and consumption rate are a function of g and does not depend on the value of $\bar{g}$. Panel A illustrates the portfolio weights of the investor. The concrete line gives the allocation to the ESG fund and the dashed line gives the allocation to the market index. Panel B shows the optimal consumption rate. Panel C to E shows the welfare of the investor in the case where $\bar{g}$ is equal to 0.25, 0.75 and 1.25 respectively. Panel F shows the green tilt of the investor, as defined in section 3. In Panel A, the solid and dashed line represents the allocation to the ESG fund and the market index respectively. In Panel C to E, the solid line represents the case where the investor observed $\bar{g}$ and the dashed line represents the case where the investor does not observe $\bar{g}$.

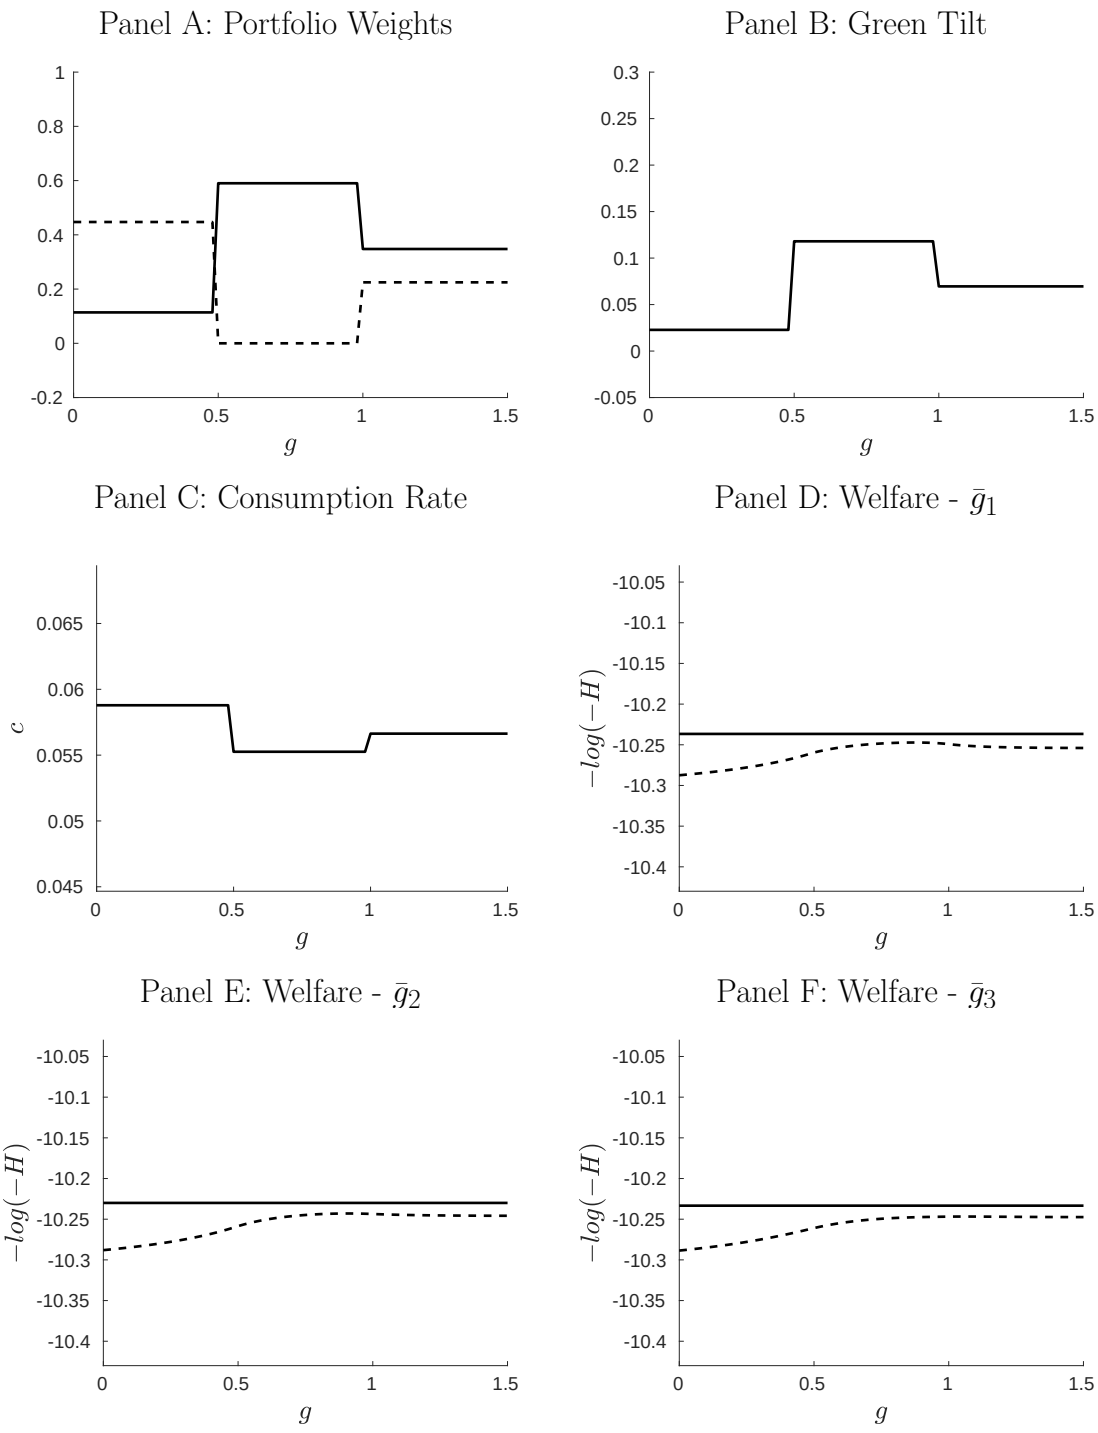
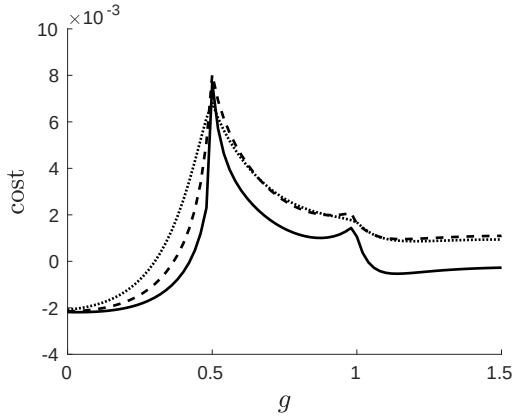


Figure 10: Cost of ESG Fund Level Uncertainty for Threshold Investor

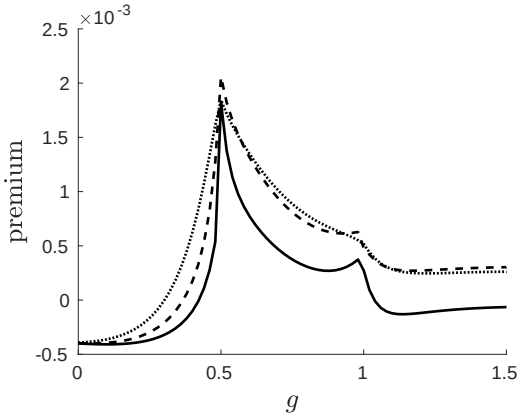
This figure decomposes the cost of ESG rating uncertainty into measurement and greenwashing uncertainty costs, for threshold investors. Panels A and B represent, respectively, the welfare cost and the return premium of measurement uncertainty. Economy 1 ($\sigma_g = 0$) is compared to Economy 0 (baseline). Panels C and D display the total welfare cost and return premium of ESG uncertainty (measurement + greenwashing). Economy 2 ($g = \bar{g}$) is compared to Economy 0. The plain, dashed and dotted line correspond to the cases where $\bar{g}$ is equal to $0.25, 0.75$ and 1.25, respectively.

Cost of measurement uncertainty: Economy 1 versus Economy 0

Panel A: Welfare cost

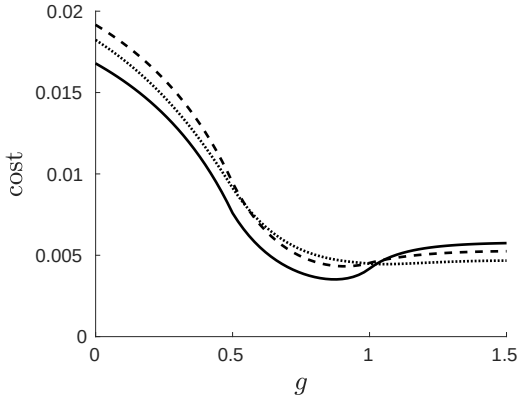


Panel B: Return premium



Cost of ESG rating uncertainty: Economy 2 versus Economy 0

Panel C: Welfare cost



Panel D: Return premium

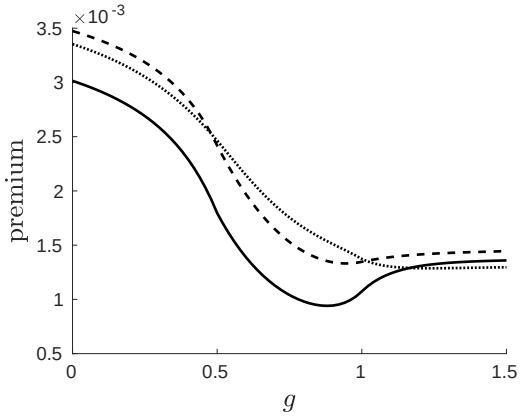


Figure 11: Sensitivity Analysis with respect to b, ESG Investor

Welfare cost of not observing the true level of greenness for an investor with utility function (13) in scenario one. The welfare cost, allocation to the ESG fund and consumption rate are computed for $b = 0.15, 0.2, 0.25$. All other parameters take their benchmark value.

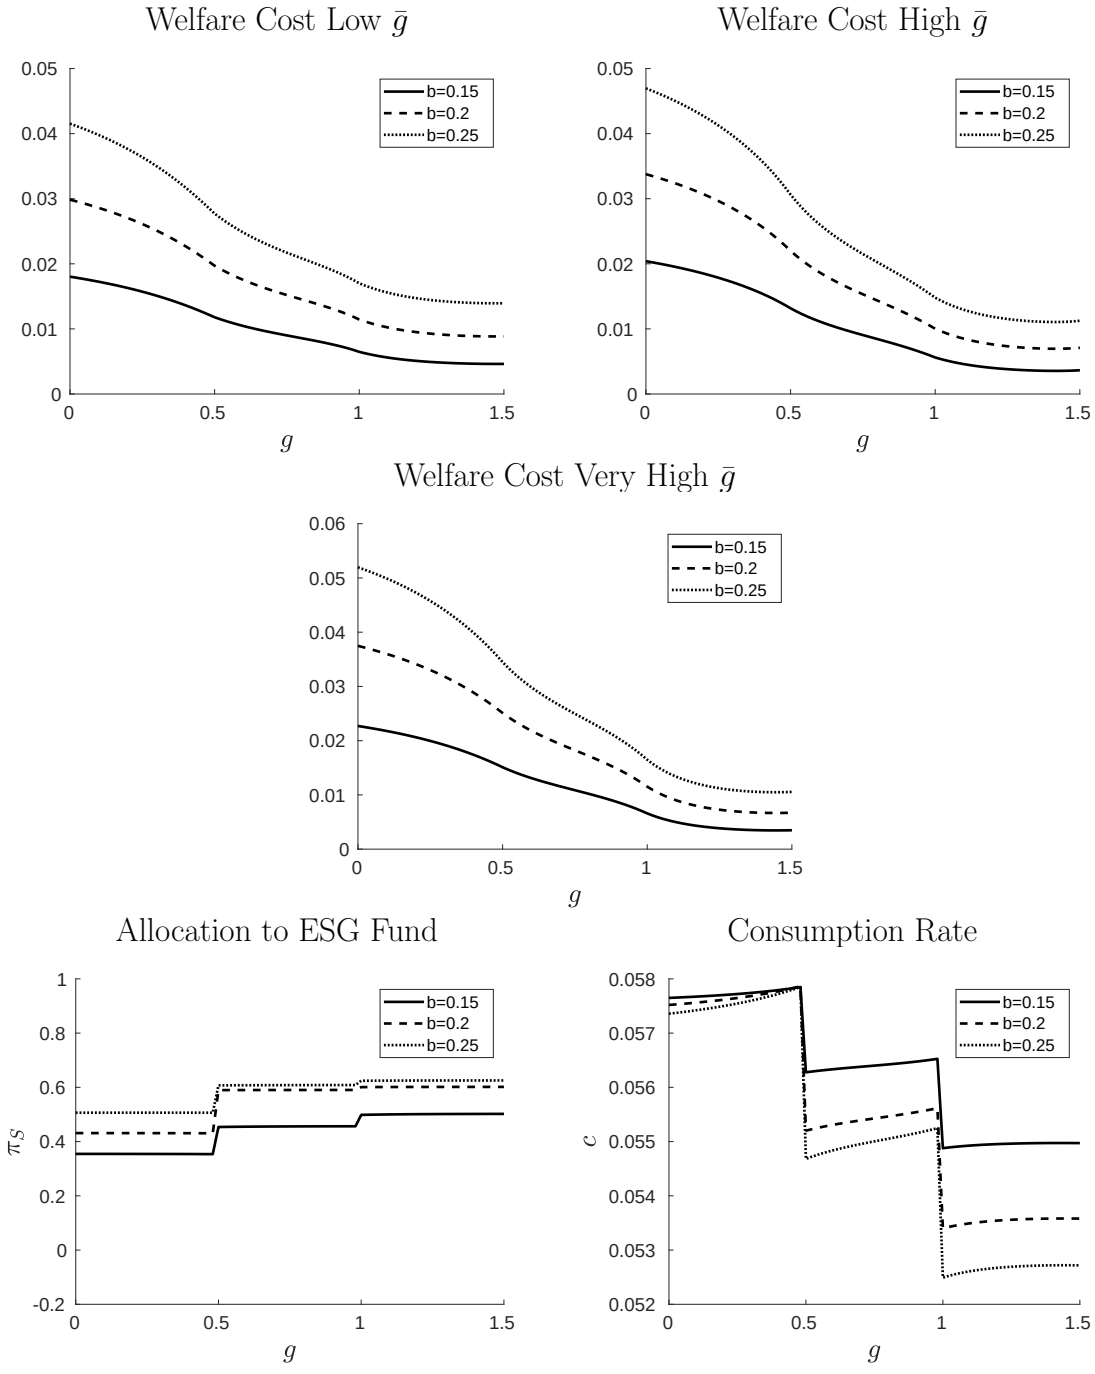


Table 1: Number of Investors in each Group

The table reports the number of investors of each category after they are grouped three equal size bins according to their AUM. There around 900 investors in each of the small, medium and large group.

	Banks	Insurers	Investment Companies	Advisors	Other	Total
Small	19	0	54	642	171	886
Medium	24	2	18	562	298	904
Large	45	8	9	344	502	908
Total	88	10	81	1548	971	2698

Table 2: AUM of Investors in each Group

The table reports the AUM of each group of investors, expressed as a percentage of the total market capitalisation.

	Banks	Insurers	Investment Companies	Advisors	Other	Total
Small	0.01	0	0.01	0.23	0.07	0.32
Medium	0.04	0	0.02	0.74	0.42	1.22
Large	13.98	2.29	1.07	19.52	61.57	98.44
Total	14.02	2.29	1.11	20.49	62.06	100

Table 3: Average Number of Incidents

The table reports the estimate of ν_1 and ν_2 from equation (18). We divide green firms into two group according to their ESG rating. The first column gives the number of firms that are moderately green and reported in each year divided by the total number of moderately green firms. The second column gives the number of very green firms that are reported in each year divided by the total number of very green firms.

Year	Estimate of ν_1	Estimate of ν_2
2007	0.0703	0.1387
2008	0.0994	0.2007
2009	0.0924	0.1861
2010	0.1345	0.2080
2011	0.1396	0.2044
2012	0.1847	0.2847
2013	0.2088	0.3029
2014	0.2319	0.3175
2015	0.2681	0.3540
2016	0.2590	0.3796
2017	0.2751	0.3796
2018	0.2550	0.3066
2019	0.3082	0.3978
2020	0.2912	0.3686
2021	0.2741	0.3905
2022	0.2902	0.3978

Table 4: Calibrated parameters

This table presents our calibrated parameters, following the calibration described in Section 5.

Parameter	Value	Definition
ESG fund dynamics		
$\bar{f}$	0.08	Long term mean of the expected return of ESG fund
θ_S	0.006	Effect of greenness on short-term expected return
σ_S	0.15	Return volatility of the ESG fund
ρ_M	0.95	Correlation between ESG stock and the market index
Dynamics of the signal		
κ_g	0.04	Mean reversion parameter of the ESG rating
σ_g	0.12	ESG rating volatility
ρ_g	0.02	Correlation between ESG rating and ESG fund return
Λ	$\begin{bmatrix} 0 & 1/2 & 0 \\ 1/4 & 0 & 1/4 \\ 0 & 1/2 & 0 \end{bmatrix}$	Transition matrix for $\bar{g}_t$
Other parameters		
μ	0.08	Expected return of the market index
σ_P	0.15	Return volatility of the market index
r	0.03	Risk free rate
γ	4	Risk aversion of the investor
b_1	0.2	Parameter of the green taste of investor one
β	0.1	Time discount parameter of the utility function
χ	0.5	Threshold of ESG signal in equation (14)

Table 5: Tilts of mutual funds

This table presents the green, brown and green minus brown (GBM) tilts of mutual funds. The tilts are calculated using definitions (2), (3) and (4). The funds labelled "ESG funds" are mutual funds with a five-globe Morningstar rating.

	# obs.	Mean	Std	Min	25th perc.	Median	75th perc.	Max
Tilts of all funds								
Green Tilt	2047	0.1403	0.1674	0.0000	0.0188	0.0770	0.2103	1.1897
Brown Tilt	2047	0.0368	0.0665	0.0000	0.0000	0.0021	0.0450	0.4675
GMB Tilt	2047	0.1035	0.2013	−0.4291	−0.0236	0.0602	0.2080	1.1897
Tilts of ESG funds								
Green Tilt	205	0.3391	0.2673	0.0000	0.1284	0.2757	0.4577	1.1897
Brown Tilt	205	0.0102	0.0252	0.0000	0.0000	0.0000	0.0006	0.1462
GMB Tilt	205	0.3289	0.2777	−0.0913	0.1263	0.2757	0.4577	1.1897

A Solving the Optimisation Problem

The expected utility of B involves two state variables, the expected return f_t and the ESG rating g_t. The case of investor A is simpler and we only need to consider f_t. This section describes the optimisation process of investor B, which can be simplified down to the optimisation process of A by replacing g_t with $\bar{g}$ in the utility function.

The transition intensities to the two states of $\bar{g}$ are equal and the dynamics of the wealth process and state variables are the same in the two states. Therefore, the two value functions are characterised by the same HJB equation with jumps (for simplicity we omit the subscript t and write $\mu_S = f - \theta_S \bar{g} - r, \mu_P = \mu - r$)

$$\begin{aligned}\beta J = \max_{\pi, C} \Big\{ & U(C,\pi) + J_W \left[(r + \pi_S \mu_S + \pi_P \mu_P)W - C\right] + \frac{1}{2} W^2 J_{WW} \left(\pi_S^2 \sigma_S^2 + \pi_P^2 \sigma_P^2 + 2\rho_M \pi_S \pi_P \sigma_S \sigma_P\right) \\ & + W J_{Wf} \pi_S \rho_f \sigma_S \sigma_f + W J_{Wg} \pi_S \rho_g \sigma_S \sigma_g + J_f \mu_f + \frac{1}{2} J_{ff} \sigma_f^2 + J_g \mu_g + \frac{1}{2} J_{gg} \sigma_g^2 + \lambda (J^\Delta - J) \Big\}\end{aligned} \tag{A.21}$$

with J^Δ denoting the value function of the alternative state that $\bar{g}$ jumps to. μ_f and μ_g are the drift term of the process f_t and g_t respectively. The optimal consumption rate and portfolio weight satisfy the first order conditions

$$\begin{aligned} 0 &= \partial U / \partial C - J_W \\ 0 &= \frac{\partial U}{\partial \pi_S} + W J_W \mu_S + W^2 J_{WW} \sigma_S^2 \left(\pi_S + \rho_M \frac{\sigma_P}{\sigma_S} \pi_P\right) + \rho_f \sigma_S \sigma_f W J_{Wf} + \rho_g \sigma_S \sigma_g W J_{Wg} \\ 0 &= W J_W \mu_P + W^2 J_{WW} \sigma_P^2 \left(\pi_P + \rho_M \frac{\sigma_S}{\sigma_P} \pi_S\right) \end{aligned} \tag{A.22}$$

To solve the HJB equation, we make the Ansatz for the value function

$$J(W, f, g) = W^{1-\gamma} H(f, g) \tag{A.23}$$

for some twice differentiable function H. The first order conditions implies that the optimal controls are given by

$$c^* = \left(\frac{(1-\gamma)H}{K}\right)^{-1/\gamma} \tag{A.24}$$

$$\pi_S^* = \frac{1}{\gamma} \left(\frac{\mu_S}{\sigma_S^2} + \frac{\rho_f \sigma_f H_f}{\sigma_S H} + \frac{\rho_g \sigma_g H_g}{\sigma_S H} + \frac{(c^*)^{1-\gamma}}{\sigma_S^2 (1-\gamma) H} (1 + b\pi_S^* g)^{-\gamma} b g\right) - \rho_M \frac{\sigma_P}{\sigma_S} \pi_P^* \tag{A.25}$$

$$\pi_P^* = \frac{1}{\gamma} \frac{\mu_P}{\sigma_P^2} - \rho_M \frac{\sigma_S}{\sigma_P} \pi_S^* \tag{A.26}$$

with $c = C/W$ denoting the consumption rate and $K = (1 + b\pi_S g)^{1-\gamma}$. Notice that there is not a closed form expression for $\pi = (\pi_S, \pi_P)$. Due to the existence of the ESG term in

the utility function, the expression of π becomes more complicated and we cannot find an analytical solution to the HJB equation (A.21). Therefore, we need to use Markov Chain Approximation method to find the solution numerically.

Equation (A.21) can be simplified into

$$0 = \frac{c^{1-\gamma}}{1-\gamma}(1+b\pi g)^{1-\gamma} + \lambda H^{\Delta} + A_0 H + A_f H_f + A_g H_g + \frac{\sigma_f^2 H_{ff} + \sigma_g^2 H_{gg}}{2} \tag{A.27}$$

with

$$\begin{aligned}
A_0 &= -\beta + (1-\gamma)(r + \pi_S\mu_S + \pi_P\mu_P - c) - \frac{\gamma}{2}(1-\gamma)(\pi_S^2\sigma_S^2 + \pi_P^2\sigma_P^2 + 2\rho_M\pi_S\pi_P\sigma_S\sigma_P) - \lambda \\
A_f &= (1-\gamma)\pi_S\rho_f\sigma_S\sigma_f + \mu_f \\
A_g &= (1-\gamma)\pi_S\rho_g\sigma_S\sigma_g + \mu_g
\end{aligned}$$

We can discretise this equation by following the process given in **?**. Consider the two dimensional grid of the values for f and g and let $z_{i,j} = (f_i, g_j)$ and $D_{i,j} = H(f_i, g_j)$. Let h_f, h_g denote the step size of the change in f, g respectively. The partial derivatives of the value function H is approximated by

$$\begin{aligned}
&\frac{\partial H(z_{i,j})}{\partial f} = \frac{D_{i+1,j} - D_{i,j}}{h_f} \text{ if } A_f \geq 0, \quad \frac{\partial H(z_{i,j})}{\partial f} = \frac{D_{i,j} - D_{i-1,j}}{h_f} \text{ if } A_f < 0, \\
&\frac{\partial H(z_{i,j})}{\partial g} = \frac{D_{i,j+1} - D_{i,j}}{h_g} \text{ if } A_g \geq 0, \quad \frac{\partial H(z_{i,j})}{\partial g} = \frac{D_{i,j} - D_{i,j-1}}{h_g} \text{ if } A_g < 0, \\
&\frac{\partial^2 H(z_{i,j})}{\partial f^2} = \frac{D_{i+1,j} - 2D_{i,j} + D_{i-1,j}}{h_f^2}, \quad \frac{\partial^2 H(z_{i,j})}{\partial g^2} = \frac{D_{i,j+1} - 2D_{i,j} + D_{i,j-1}}{h_g^2}
\end{aligned}$$

Plugging the approximations above into the HJB equation gives, for all i, j (We write $X^+ = \max(X, 0)$, $X^- = \max(-X, 0)$),

$$\begin{aligned}
0 = {} & \frac{c^{1-\gamma}}{1-\gamma}(1+b\pi g)^{1-\gamma} + \lambda D_{i,j}^{\Delta} + A_0 D_{i,j} + A_f^+ \frac{D_{i+1,j} - D_{i,j}}{h_f} - A_f^- \frac{D_{i,j} - D_{i-1,j}}{h_f} \\
& + A_g^+ \frac{D_{i,j+1} - D_{i,j}}{h_g} - A_g^- \frac{D_{i,j} - D_{i,j-1}}{h_g} + \frac{\sigma_f^2}{2}\frac{D_{i+1,j} - 2D_{i,j} + D_{i-1,j}}{h_f^2} \\
& + \frac{\sigma_g^2}{2}\frac{D_{i,j+1} - 2D_{i,j} + D_{i,j-1}}{h_g^2}
\end{aligned} \tag{A.28}$$

Rearranging gives

$$\begin{aligned}
D_{i,j} = {} & \frac{c^{1-\gamma}(1+b\pi g)^{1-\gamma}}{(1-\gamma)\delta(z_{i,j}, c, \pi)} + \frac{\lambda D_{i,j}^{\Delta}}{\delta(z_{i,j}, c, \pi)} + p(z_{i,j}, z_{i+1,j}|c, \pi)D_{i+1,j} \\
& + p(z_{i,j}, z_{i-1,j}|c, \pi)D_{i-1,j} + p(z_{i,j}, z_{i,j+1}|c, \pi)D_{i,j+1} + p(z_{i,j}, z_{i,j-1}|c, \pi)D_{i,j-1}
\end{aligned} \tag{A.29}$$

The equation above characterises a Markov chain with the transition probabilities given by

(We write $X^+ = \max(X, 0)$, $X^- = \max(-X, 0)$)

$$
\begin{aligned}
\delta(z_{i,j}, c, \pi) &= -A_0 + \frac{|A_f|}{h_f} + \frac{|A_g|}{h_g} + \frac{\sigma_f^2}{h_f^2} + \frac{\sigma_g^2}{h_g^2}, \\
p(z_{i,j}, z_{i+1,j}|c, \pi) &= \frac{1}{\delta(z_{i,j}, c, \pi)} \left(\frac{A_f^+}{h_f} + \frac{\sigma_f^2}{2h_f^2} \right), \\
p(z_{i,j}, z_{i-1,j}|c, \pi) &= \frac{1}{\delta(z_{i,j}, c, \pi)} \left(\frac{A_f^-}{h_f} + \frac{\sigma_f^2}{2h_f^2} \right), \\
p(z_{i,j}, z_{i,j+1}|c, \pi) &= \frac{1}{\delta(z_{i,j}, c, \pi)} \left(\frac{A_g^+}{h_g} + \frac{\sigma_g^2}{2h_g^2} \right), \\
p(z_{i,j}, z_{i,j-1}|c, \pi) &= \frac{1}{\delta(z_{i,j}, c, \pi)} \left(\frac{A_g^-}{h_g} + \frac{\sigma_g^2}{2h_g^2} \right)
\end{aligned}
\tag{A.30}
$$

A.1 Value Function Ansatz

It is convenient to calculate that, in the power utility case,

$$
\begin{aligned}
&WJ_W = (1-\gamma)J \\
&W^2 J_{WW} = -\gamma(1-\gamma)J \\
&WJ_{Wf} = (1-\gamma)\frac{H_f}{H}J \\
&WJ_{Wg} = (1-\gamma)\frac{H_g}{H}J \\
&J_f = \frac{H_f}{H}J \\
&J_g = \frac{H_g}{H}J \\
&J_{ff} = \frac{H_{ff}}{H}J \\
&J_{gg} = \frac{H_{gg}}{H}J
\end{aligned}
$$

B Find the expected utility with sub-optimal controls

Given an asset allocation strategy, characterised by $\pi(g_t, \bar{g}_t), c(g_t, \bar{g}_t)$, the cumulative expected utility is

$$
J(W_t, g_t, \bar{g}_t) = \mathbb{E}_t \left(\int_t^\infty e^{-\beta(u-t)} U(C_u, \pi_u, g_u) du \right)
$$

Here we assume J only depends on state variables and not t since our problem is for an infinite time horizon. We denote the Dynkin operator by $\mathcal{L}$, with

$$\begin{aligned}\mathcal{L}_t J = &\frac{\partial J_t}{\partial W} W_t(r - c_t + \pi_t \mu(t)) + \frac{1}{2}\frac{\partial^2 J_t}{\partial W^2} W_t^2 \pi_t^2 \sigma^2 + \frac{\partial J_t}{\partial g}\mu_g(t) + \frac{1}{2}\frac{\partial^2 J_t}{\partial g^2}\sigma_g^2 \\ &+ \frac{\partial^2 J_t}{\partial W \partial g} W_t \pi_t \rho_g \sigma \sigma_g\end{aligned}$$

Then we define

$$Y(s) = \int_t^{t+s} e^{-\beta(u-t)} U(C_u, \pi_u, g_u) du + e^{-\beta s} J(W_{t+s}, g_{t+s}, \bar{g}_{t+s})$$

Apply Ito's formula to $Y(s)$ gives

$$\begin{aligned}Y_s - Y_0 &= -\beta e^{-\beta s} J(W_{t+s}, g_{t+s}, \bar{g}_{t+s}) ds + e^{-\beta s} dJ(W_{t+s}, g_{t+s}, \bar{g}_{t+s}) + e^{-\beta s} U(C_{t+s}, g_{t+s}, \bar{g}_{t+s}) ds \\ &= e^{-\beta s}\Bigg[- \beta J(W_{t+s}, g_{t+s}, \bar{g}_{t+s}) + \int_0^s \mathcal{L}_{t+s} J + U(C_{t+s}, g_{t+s}, \bar{g}_{t+s}) ds + \frac{\partial J_{t+s}}{\partial g}\sigma_g dZ_t^g \\ &+ \frac{\partial J_{t+s}}{\partial W} W_{t+s} \pi_{t+s} \sigma_S dZ_t^s \Bigg] + \sum_{0<\tau<s : \bar{g}_{t+\tau} \neq \bar{g}_{t+\tau}^-} e^{-\beta\tau}\left(J_{t+\tau} - J_{t+\tau}^-\right)\end{aligned}$$

Recall that we $\bar{g}$ jumps to the alternative state following a Poisson process with intensity λ. Let J^Δ denote the the value function in the alternative state of $\bar{g}$, we have

$$\lim_{s\to 0} \mathbb{E}_t \left(\sum_{0<\tau<s : \bar{g}_{t+\tau} \neq \bar{g}_{t+\tau}^-} e^{-\beta\tau} \frac{J_{t+\tau} - J_{t+\tau}^-}{s} \right) = \lambda(J_t^\Delta - J_t)$$

Then notice that

$$\mathbb{E}_t(Y_s) = Y_0$$

so

$$0 = \lim_{s\to 0} \mathbb{E}_t \left(\frac{Y_s - Y_0}{s} \right) = -\beta J_t + U(C_t, g_t, \bar{g}_t) + \mathcal{L}_t J + \lambda(J_t^\Delta - J_t)$$

which is the differential equation that characterises the sub-optimal expected utility. This equation coincides with the HJB equation with optimal controls. This means we can just plug the sub-optimal controls into the algorithm for numerically solving the HJB equation to compute the expected utility in the sub-optimal case.

C Fitting parameters for the model

We follow the procedure described by Holmes (2003) to calibrate the parameters in the dynamics of stock processes. First, we define a function based on Kalman filter method that estimates the expected return from realised stock returns. Then we apply the maximum likelihood estimate to the conditional probability distribution of stock return predicted by the Kalman filter to find $\kappa_f, \sigma_f, \rho_f$.

In particular, we discretise the process of stock price and expected return and get

$$\log(S_{t+1}) - \log(S_t) = f_t\Delta_t + \sigma_S\sqrt{\Delta_t}\varepsilon_t \tag{C.31}$$

$$f_t = (1-\kappa_f\Delta_t)f_{t-1} + \kappa_f\Delta_t\bar{f} + \sigma_f\sqrt{\Delta_t}\iota_{t-1} \tag{C.32}$$

where ε_t, ι_t are i.i.d standard normal. Then by defining $y_t = \log(S_{t+1}) - \log(S_t)$ and $x_t = f_t\Delta_t$ we have

$$y_t = x_t + \nu_t \tag{C.33}$$

$$x_t = Ax_{t-1} + B + \omega_{t-1} \tag{C.34}$$

with $A = 1-\kappa_f\Delta_t$, $B = \bar{f}\kappa_f\Delta_t^2$, $\nu_t \overset{i.i.d}{\sim} \mathcal{N}(0, \sigma_S\sqrt{\Delta_t})$, $\omega_t \overset{i.i.d}{\sim} \mathcal{N}(0, \sigma_f\Delta_t^{3/2})$. To simplify the notation, we define $Q = \sigma_f\Delta_t^{3/2}$ and $R = \sigma_S\sqrt{\Delta_t}$. To apply the standard Kalman filter, we need $\mathrm{Cov}(\nu_t, \omega_s) = 0$ for all t, s. However, the correlation between ν_t, ω_t in our case is ρ_f, which is allowed to be nonzero. Therefore, we need to decorrelate x_t, y_t by following Chang (2014) and rewrite the state system as

$$y_t = x_t + \nu_t \tag{C.35}$$

$$x_t = A^*x_{t-1} + u_{t-1} + \omega^*_{t-1} \tag{C.36}$$

with

$$A^* = 1 - \kappa_f\Delta_t - \rho_f\sqrt{Q/R} \tag{C.37}$$

$$u_{t-1} = \bar{f}\kappa_f\Delta_t^2 + y_{t-1}\rho_f\sqrt{Q/R} \tag{C.38}$$

$$\omega^*_{t-1} = \omega_{t-1} - \nu_{t-1}\rho_f\sqrt{Q/R} \tag{C.39}$$

$$\mathrm{Var}(\omega^*_t) = (1-\rho_f^2)Q \tag{C.40}$$

Since $\mathrm{Cov}(\omega^*_t, \nu_t) = 0$, we can apply the standard Kalman filter to (C.35) and (C.36).