Green Investing Under Limited Attention

M. Donadelli, A. Mazzoran, R. Kizys*

January 29, 2026

Abstract

We develop a two-period rationally inattentive portfolio-choice model to explain the slow transition of investor portfolios towards sustainability. When information-processing capacity κ is limited, and signals about returns and environmental damage are noisy, investors remain imperfectly informed about brown transition risks and green opportunities. This leads to persistent underweighting of green assets, overexposure to brown assets, sluggish price reactions to climate-policy shocks, and high cross-sectional dispersion in portfolio sustainability even among investors with homogeneous preferences. We show that meaningful brown divestment and green overweighting arise only once information-processing capacity exceeds a sharp κ-threshold. This threshold declines with stronger environmental preferences, improved signal quality, fixed consumption, or a convex environmental penalty. Extensions with quadratic disutility from brown exposure amplify these effects: easing the attention constraint even slightly triggers rapid, near-complete reallocation from brown to green assets. Overall, the model reconciles the widespread gap between investors' stated sustainability preferences and observed portfolio choices, rationalizes delayed price responses to green-policy news, and provides micro-foundations for the sluggish greening of the financial system.

1 Introduction

The rapid shift toward sustainable finance has fundamentally altered how investors allocate capital between green and brown stocks. Green stocks – firms that integrate environmental considerations into their operations – benefit from growing investor preferences for environmental sustainability, regulatory support, and consumer demand for ethical products. In contrast, brown stocks – firms operating in industries that contribute significantly to environmental degradation, such as fossil fuels, heavy manufacturing, and mining – face increasing headwinds from climate policies and environmental regulations. As global sustainability standards evolve, the financial performance of these asset classes remains subject to substantial uncertainty, creating a pressing need to understand the mechanisms driving their valuation.

*Michael Donadelli is at the Department of Economics and Management of the University of Brescia (michael.donadelli@unibs.it). Renatas Kizys is affiliated with the University of Southampton Business School (R.Kizys@soton.ac.uk; renatas.kizys@googlemail.com). Andrea Mazzoran is at the Department of Quantitative Finance of the University of Freiburg (andrea.mazzoran@finance.uni-freiburg.de).

Acknowledgements: We thank Carmen Marchiori, Jerome Detemple, Luca Lochi, Abhishek Kumar, and participants in several seminars and conferences for valuable comments and suggestions. The authors received no external financial support for this research. The authors declare that they have no financial relationships or other potential conflicts of interest relevant to this work. For guidance on disclosures, we have consulted the NBER research dissemination disclosure policy.

The foundational work of Pastor et al. (2021) and Pedersen et al. (2021) provides a theoretical cornerstone for understanding how investors price green and brown assets. Their equilibrium models show that when investors exhibit preferences for environmental sustainability, green assets command lower expected returns, reflecting a *green premium*, which arises from non-pecuniary utility rather than superior financial performance. Yet, these static frameworks struggle to reconcile this prediction with the substantial realized returns that green investments have generated in recent years – a period characterized by rapid changes in investor sentiment, evolving climate and environmental policies, and heightened awareness of environmental risks. Our model helps explain this puzzle by showing that when investors face information-processing limits, green assets can remain underweighted for prolonged periods, so that positive environmental shocks subsequently generate excess realized returns as attention gradually shifts toward them.

Dynamic models, such as Avramov et al. (2025), account for time-varying investor preferences towards green assets by treating shifts in aggregate demand for such investments as priced sources of risk. In their framework, unexpected increases in investor appetite for green assets can generate substantial positive deviations between realized and expected returns on green assets. Although these models can reproduce periods of strong green performance, they implicitly assume that investors process climate- and environmental-related information efficiently. In contrast, empirical evidence indicates that climate-related news introduces dynamic, time-varying risks that require active hedging (Engle et al., 2020), and that investors gradually reward green funds (Hartzmark and Sussman, 2019) while pricing firms' carbon exposure (Bolton and Kacperczyk, 2021). Across these studies, markets adjust only gradually to climate- and environmental-related news. Thus, although both theoretical and empirical strands agree that investor preferences for green assets affect asset prices, they diverge on the mechanism and timing of price adjustment. This divergence motivates our central question: why does climate- and environmental-related information diffuse so slowly into asset prices?

However, this dynamic perspective raises a deeper question: why do shifts in investor demand for green or climate-friendly assets exert such persistent and pronounced effects on prices, rather than being swiftly incorporated into market valuations? In an efficient market where climate- and environmental-related information is immediately processed by all investors, such shifts should be rapidly arbitraged away. The persistence of these return effects suggests the presence of frictions – whether informational, behavioral, or institutional – that slow the transmission of climate- and environmental-related information into prices. Yet, the existing literature, focused primarily on equilibrium outcomes and aggregate empirical patterns, offers little in the way of a micro-founded account of how and why investor portfolios adjust sluggishly to evolving climate- and environmental-related signals.

We argue that the missing mechanism underlying the sluggish adjustment of asset prices to sustainability information is rational inattention (RI) (Sims, 2003). Investors face inherent cognitive and informational constraints when processing complex and noisy signals about firms' environmental performance, regulatory developments, and transition risks. While dynamic asset-pricing frameworks such as Avramov et al. (2025) capture the consequences of slow-moving demand for green investments, they remain silent on its origins. To bridge this gap, and follow-

ing the spirit of Batchuluun et al. (2019), we introduce a RI constraint into a portfolio choice setting, providing a micro-founded explanation for why investors respond gradually to climate- and environmental-related news. Building on the application of RI in financial markets by Peng and Xiong (2006) and Van Nieuwerburgh and Veldkamp (2010), we develop a tractable two-period model that isolates the role of information-processing capacity (κ) in shaping investors' optimal allocation among green, brown, and the risk-free assets.

Our theoretical model and numerical solutions yield a consistent set of results that underscore the significant role of limited attention. We find that investors with low information-processing capacity exhibit slower responses to climate- and environmental-related news, leading to a persistent underweighting of green assets. An increase in capacity κ systematically boosts allocations to green investments, while higher risk aversion paradoxically shifts portfolios towards brown assets, which offer a risk premium. Moreover, the quality of information itself is decisive: noisier signals for green stocks suppress their allocation, whereas noisier signals for brown stocks can indirectly promote green investment. These cognitive frictions generate delayed market adjustments and persistent pricing inefficiencies, and help explain why climate- and environmental-related shocks have such prolonged effects on returns.

Overall, we demonstrate that inefficiencies in green investing arise not only from regulatory uncertainty and behavioral biases but also from fundamental cognitive constraints. This research seeks to answer the following key questions: First, how does limited information-processing capacity influence the allocation of wealth between green and brown stocks? Second, what role does attention allocation play in shaping the speed of market adjustment to climate- and environmental-related shocks? Third, under what conditions do investors underreact to new climate- and environmental-related regulations?

This analysis allows us to make the following key contributions to the literature on green finance and information frictions. First, we propose a novel integration of the RI framework into a green-versus-brown portfolio choice model – an approach that, to the best of our knowledge, has no precedent in the existing literature. This integration provides a micro-founded explanation for the observed sluggish adjustment of asset prices and investor portfolios to environmental news, enhancing and extending prior explanations that predominantly rely on regulatory uncertainty or behavioral biases. Second, we demonstrate that strong environmental preferences (ϕ) are a necessary but not sufficient condition for green investing. Our model shows that these preferences translate into higher green allocations only when complemented by sufficient information-processing capacity (κ), highlighting a previously overlooked cognitive bottleneck in green finance. Third, we identify a paradoxical driver of green investment: higher signal noise for brown assets. When information about brown stocks becomes noisier, investors rationally shift attention and capital towards green alternatives, even without changes in their environmental preferences. Fourth, the model predicts that the temporary mispricing created by rationally inattentive investors generates distinct arbitrage opportunities for more attentive investors, who can adjust their portfolios earlier in response to environmental shocks. Fifth, through numerical solutions and sensitivity analysis, we quantify how portfolio choices and adjustment speeds are shaped by the interplay of risk aversion (γ), environmental concern (ϕ), signal noise (σ^2), and information capacity (κ), which translates into testable empirical predictions.

The remainder of this paper is structured as follows. Section 2 positions our study within the relevant literature on green finance and RI. Section 3 introduces the two-period portfolio choice model. This model features investor preferences for both financial returns and environmental impact, and incorporates the RI constraint. Section 4 derives the optimization problem and the resulting first-order conditions. It also characterizes the mechanics of optimal attention allocation. Section 5 presents numerical results, including sensitivity analyses, which quantify how portfolio choices respond to key parameters such as information capacity, risk aversion, and signal noise. Finally, Section 6 concludes.

2 Literature review

The growing prominence of sustainable finance has catalyzed a rich body of research on how ESG factors influence investor behavior and asset prices. Concurrently, a separate and well-established literature explores the profound implications of information frictions for economic and financial decision-making. This review seeks to reconcile these two strands, arguing that the theory of RI provides a micro-founded explanation for central puzzles in sustainable investing.

A central premise in the sustainable finance literature is that investors derive utility not solely from financial returns, but also from non-pecuniary benefits Seminal work by Pastor et al. (2021) formalizes this idea in an equilibrium setting, which demonstrates that investors' preferences for sustainable assets can lower their expected returns, creating a "greenium". This aligns with empirical evidence that a significant segment of the investor population is willing to sacrifice financial performance for social or environmental objectives (Krueger et al., 2020; Hartzmark and Sussman, 2019). Furthermore, assets perceived as "brown" or carrying high climate risk have been shown to exhibit different risk-return profiles. These assets often require a financial premium to compensate for their exposure to transition risks, such as future carbon taxes or regulatory shifts (Bolton and Kacperczyk, 2021; Görgen et al., 2020). These studies establish a clear theoretical and empirical basis for distinguishing between green and brown assets and for incorporating a direct environmental impact term into the investor's utility function.

However, the literature remains largely silent on the informational channels and cognitive processes through which investors operationalize these preferences. ESG-related information is often complex, forward-looking, and non-standard, pertaining to long-term risks like climate change or social governance. A wealth of evidence from behavioral finance documents systematic underreaction to such information. For instance, studies show that stock prices incorporate news about firm-specific ESG performance only gradually (Dimson et al., 2015), and markets are slow to price in long-term climate risks (Giglio et al., 2021). Traditional explanations for this underreaction often point to behavioral biases or regulatory uncertainty. Yet, a more foundational explanation lies in the theory of RI, pioneered by Sims (2003).

The concept of RI represents a paradigm shift in the way economists model decision-making under uncertainty. Taking its origins from the seminal work of Sims (2003), RI theory posits that economic agents face fundamental constraints in their capacity to process information. Unlike traditional models where information is either fully available or completely unavailable, the RI framework recognizes that agents can process information but might optimally choose to

do so imperfectly. Building on information theory, RI is formalized as a constraint on mutual information between the state variables and the agent's signals. This is typically expressed as $I(\theta;s) \leq \kappa$, where θ represents the state variables, s denotes the signals observed by the agent, and κ captures the agent's information-processing capacity. This constraint captures the trade-off between signal precision and cognitive effort.

The early applications of RI primarily focused on macroeconomic phenomena, particularly price stickiness. Mackowiak and Wiederholt (2009) demonstrated how rationally inattentive price-setters could generate persistent nominal rigidities without requiring menu costs or other frictions. Their model showed that when firms face capacity constraints in processing macroeconomic information, they optimally adjust prices infrequently and by small amounts, which generates realistic inflation dynamics. This line of research provided an important alternative to traditional sticky-price models by emphasizing the cognitive constraints rather than physical adjustment costs. The RI framework successfully explained why firms might not immediately respond to monetary policy shocks or other aggregate disturbances, even when such responses would be beneficial under full information.

In financial economics, RI has proven particularly powerful in explaining asset pricing puzzles and investor behaviors. Peng and Xiong (2006) showed how inattention to market-wide information can generate momentum and reversal patterns in stock returns. Their model demonstrated that when investors have limited capacity to process market information, they tend to underreact to macroeconomic news initially but overreact in the long run as information gradually diffuses. A key insight from this literature is that rational inattention naturally generates category-learning behavior, where investors allocate their attention broadly to market or sector-level information rather than processing firm-specific details. This explains why correlated price movements occur within asset classes and why investors might appear to neglect idiosyncratic information.

Van Nieuwerburgh and Veldkamp (2010) extended the RI framework to portfolio choice, which shows how attention constraints lead to suboptimal diversification and slow portfolio rebalancing. Their model predicts that investors concentrate their attention on assets they already hold heavily, creating inertia in portfolio adjustments; a finding with direct implications for understanding the slow diffusion of ESG information into investment portfolios.

The allocation of attention across different information sources represents a central theme in RI research. Kacperczyk et al. (2016) developed a model where investors optimally allocate their limited attention between learning about asset fundamentals and monitoring portfolio performance. They showed that this attention allocation decision has significant implications for trading behavior, portfolio concentration, and ultimately, investment performance. Their empirical tests revealed that more sophisticated investors (those with higher κ) exhibit better attention allocation and, consequently, achieve superior performance. This heterogeneity in information-processing capacity provides an explanation for why some investors can exploit market inefficiencies, while others persistently underperform.

The implications of RI for market efficiency have been extensively explored. Hirshleifer et al. (2011) demonstrated how inattention can create predictable return patterns around corporate earnings announcements and other information events. Their model shows that even when

information is publicly available, the gradual processing of this information by inattentive investors generates persistent mispricing. Importantly, the RI framework distinguishes itself from behavioral models by maintaining the assumption of rationality within cognitive constraints. Investors in RI models do not make systematic cognitive errors. Rather, they optimally allocate their scarce attention resources. This distinction is crucial for understanding why mispricing patterns can persist even in the presence of sophisticated arbitrageurs, as attention allocation itself becomes a strategic decision.

In parallel, the literature on sustainable finance and environmental, social, and governance (ESG) investing has highlighted how preferences and information jointly drive portfolio decisions. Pastor et al. (2021) developed an equilibrium model of sustainable investing, which demonstrates that investors derive utility from holding green assets even at the cost of lower expected returns. Empirical studies, such as Krueger et al. (2020), confirm that institutional investors acknowledge climate risk but face challenges in processing ESG-related data due to measurement inconsistencies and disclosure heterogeneity. Avramov et al. (2025) documented that markets underreact to ESG demand shocks – where green assets earn higher realized returns than initially anticipated – suggesting that information frictions play a central role in the dynamics of sustainable asset pricing.

Overall, our work provides a novel contribution to the literature by integrating the RI framework directly into a sustainable portfolio choice model, thereby offering a micro-founded explanation for the observed underreaction to green-related information that moves beyond existing explanations rooted in regulatory uncertainty or behavioral biases alone. While prior work has established RI as a key driver of inertia in macroeconomic settings and traditional financial markets (Sims, 2003; Mackowiak and Wiederholt, 2009) and has modeled its implications for asset pricing (Peng and Xiong, 2006) and portfolio allocation (Van Nieuwerburgh and Veldkamp, 2010), our model is the first to explicitly formalize how a finite information-processing capacity constrains an investor's ability to process noisy signals on both the financial and environmental attributes of green and brown assets simultaneously.

More broadly, our research yields novel theoretical predictions about investor behavior and asset pricing. Specifically, we find that strong environmental preferences only translate into higher green allocations when complemented by sufficient attention capacity. We also find that a higher level of signal noise for brown assets can paradoxically boost green investment. Our theoretical model explains how these cognitive frictions can lead to persistent mispricing, delayed portfolio adjustment, and create arbitrage opportunities for attentive investors.

3 Model Setup

3.1 Investor Preferences

The investor's utility function accounts for both financial wealth and environmental concerns, which is essential given the increasing role of ESG preferences in asset allocation (Pastor et al., 2021). We assume an iso-elastic (CRRA) utility function over consumption c and a linear function in environmental impact E

$$u(c,E) = \frac{c^{1-\gamma}}{1-\gamma} - \phi E, \tag{1}$$

In Equation 1, $\gamma > 0$, $\gamma \neq 1$, is the coefficient of relative risk aversion, $\phi \geq 0$ is the weight on environmental impact, and $E = \alpha_b E_b$ is the environmental impact of the portfolio, where α_b is the fraction of wealth invested in brown stocks and $E_b > 0$ denotes the environmental damage per unit of brown stock.[1] This specification aligns with van der Ploeg (2025), who similarly assumes linear disutility from environmental degradation, and is further consistent with empirical evidence showing that investors incorporate non-financial factors into their decision-making (Krueger et al., 2020). The model allows for heterogeneous environmental preferences. Indeed, some investors may be willing to sacrifice returns to hold greener portfolios, while others may prioritize financial performance over sustainability. Investors maximize lifetime utility over two periods, in which they make decisions in the first period based on expectations of future outcomes in the second period

$$\mathbb{E}[u(c_1,E_1) + \beta u(c_2,E_2)], \tag{2}$$

where $\beta \in (0,1)$ is the discount factor. The problem consists of two key stages:

Period 1. At the beginning of the first period, the investor receives a stochastic income e_1 and chooses consumption c_1. The remaining wealth is allocated across three asset classes: green stocks (α_g) with an expected return R_g, brown stocks (α_b) with an expected return R_b, and a risk-free asset (α_f) with a fixed return R_f.

Period 2. In the second period, the investor receives an additional stochastic income e_2 and consumes the total available wealth. The portfolio return R_p determines the accumulated capital carried forward and is given by

$$R_p = \alpha_g R_g + \alpha_b R_b + \alpha_f R_f. \tag{3}$$

The second-period consumption c_2 is then

$$c_2 = (e_1 - c_1)R_p + e_2. \tag{4}$$

This formulation captures the investor's intertemporal trade-off between immediate consumption and long-term portfolio growth, subject to the uncertainty of asset returns and their environmental impact.

[1] Following Pastor et al. (2021), we model environmental preferences through a direct utility penalty $-\phi\alpha_b E_b$. This term captures a purely taste-based disutility from holding brown assets, independent of any causal effect on firm emissions. Our model therefore abstracts from feedback effects on firm behavior. This approach is consistent with the majority of ESG-portfolio models (e.g., Pastor et al., 2021; Pedersen et al., 2021), which similarly treat exclusion as preference-driven rather than impact-driven. Incorporating real effects of exclusion on firm emissions or brown returns – for example, via equilibrium firm responses – would be an important and complementary extension, but falls outside the scope of the present two-period setting.

Portfolio Choice

Investors allocate wealth between green stocks (α_g) with the expected return R_g and environmental impact $E_g = 0$, brown stocks (α_b) with the expected return R_b and environmental impact $E_b > 0$, and the risk-free asset (α_f) with the return R_f. The budget constraint is given by

$$\alpha_g + \alpha_b + \alpha_f = 1. \tag{5}$$

Information Structure and Rational Inattention

Investors receive noisy signals s_g and s_b of green and brown stocks, which are linear combinations of portfolio exposure on true future returns with noise. Formally,

$$s_g = \alpha_g R_g + \varepsilon_g, \quad s_b = \alpha_b R_b + \varepsilon_b, \tag{6}$$

where the noisy terms ε_g and ε_b are assumed to be zero-mean Gaussian with variances σ_g^2 and σ_b^2, respectively.[2] We further assume (R_g, ε_g) to be independent of (R_b, ε_b).[3] The noise in the signals affects the investor's ability to process information about the returns of green and brown stocks. For instance, higher noise reduces the investor's ability to process information about green stocks, which leads to lower allocations. As a result, they become less confident about returns on green stocks and may shift wealth to safer assets. Higher noise also reduces the investor's ability to process information about brown stocks, but the effect may be less pronounced if brown stocks are already perceived as risky. Consequently, investors may maintain or slightly reduce their allocations to brown stocks. Furthermore, as noise increases, the allocation to the risk-free asset typically increases because investors seek safer investments when information is less reliable. Due to the limited information-processing capacity κ, investors cannot fully process all the information in the signals. They allocate κ between the signals s_g and s_b to maximize their expected utility. Optimal attention allocation depends on the relative importance of the signals in determining portfolio return and environmental impact. The mutual information between the signals and the investor's decision (portfolio choice) is constrained by κ. This can be formalized as

$$I(R_g, R_b; s_g, s_b) \leq \kappa, \tag{7}$$

where $I(\cdot)$ is the mutual information. For example, if the shock increases the expected return on green stocks, investors can choose to allocate more attention to s_g to better estimate it. A

[2] In the standard rational-inattention literature, signals are often written in the reduced form $s_g = R_g + \varepsilon_g$ (see, e.g., Peng and Xiong, 2006; Van Nieuwerburgh and Veldkamp, 2010). We adopt the scaled specification $s_g = \alpha_g R_g + \varepsilon_g$ for tractability within the mutual-information framework. The scaling by α_g reflects the idea that the marginal value of information is proportional to the investor's exposure to the asset, as in entropy-based learning technologies. The economic interpretation is unchanged; investors begin with prior beliefs about returns and update them using the signals as Bayesian likelihoods. This formulation preserves the standard Bayesian updating logic while yielding a closed-form mutual-information constraint and the Gaussian assumption.

[3] The assumption of independence between green and brown asset returns is adopted primarily for analytical tractability and expositional clarity. Introducing correlation between the two return processes would significantly complicate the information structure and the resulting dynamics of optimal portfolio choice, thereby obscuring the sharp qualitative patterns that emerge under independence. We therefore leave the theoretically important case of correlated green–brown returns for future research.

shock can contain new information about the environmental impact of brown stocks (e.g. stricter regulations) or the financial performance of green stocks (e.g. renewable energy subsidies). It can be modeled as a change in the distribution of R_g and R_b, or as a change in E_g and E_b. When a shock arrives, investors update their beliefs about the returns and environmental impacts of green and brown stocks. However, due to their rational inattention, they may not fully incorporate the new information into their decision-making process.

Model Dynamics and Key Predictions

Pre-Shock Equilibrium. Prior to any shocks, investors allocate their limited attention between the signals s_g and s_b, which in turn leads to their portfolio choices (α_g, α_b). Investors with different information-processing capacity κ may exhibit different portfolio choices, where some of them hold more green stocks than others. Investors with low κ may exhibit a more precautionary behavior, adjusting their portfolios more conservatively, reflecting greater inertia in response to uncertainty.

Post-Shock Adjustment. Following a shock, investors reallocate attention in response to the new information. When the shock favors green stocks, investors optimally increase attention to s_g, leading to a higher portfolio share α_g. However, due to limited information-processing capacity, the adjustment is typically incomplete, and portfolios may not fully reflect the new information. This leads to a gradual transition towards a new equilibrium.

Long-Run Implications. Over time, as investors continue to process information, portfolio allocations gradually converge toward the optimal allocation implied by the new information. The speed of this adjustment is driven by the investor's information-processing capacity κ.

4 Optimization Problem

The mutual information constraint captures the investor's limited information-processing capacity by quantifying the reduction in uncertainty about the returns R_g and R_b after observing the signals s_g and s_b. The mutual information $I(R_g, R_b; s_g, s_b)$ is defined as

$$I(R_g, R_b; s_g, s_b) = H(R_g, R_b) - H(R_g, R_b | s_g, s_b), \tag{8}$$

where $H(R_g, R_b)$ is the joint entropy of the returns, and $H(R_g, R_b | s_g, s_b)$ is the conditional entropy of the returns given the signals. Under the independence assumptions introduced above and the joint Gaussianity of returns and signals, the mutual information decomposes additively and can be expressed in terms of variances

$$I(R_g, R_b; s_g, s_b) = \frac{1}{2} \log \left(\frac{\mathrm{Var}(R_g)}{\mathrm{Var}(R_g | s_g)} \right) + \frac{1}{2} \log \left(\frac{\mathrm{Var}(R_b)}{\mathrm{Var}(R_b | s_b)} \right), \tag{9}$$

where $\mathrm{Var}(R_g)$ is the unconditional variance of the return on green stocks, $\mathrm{Var}(R_g | s_g)$ is the conditional variance of R_g given the signal s_g, $\mathrm{Var}(R_b)$ is the unconditional variance of the return on brown stocks, $\mathrm{Var}(R_b | s_b)$ is the conditional variance of R_b given the signal s_b. The mutual

information constraint limits the investor's ability to reduce uncertainty about the returns R_g and R_b by observing the signals s_g and s_b. In this expression, the term $\frac{1}{2}\log\left(\frac{\text{Var}(R_g)}{\text{Var}(R_g|s_g)}\right)$ captures the reduction in uncertainty about green return R_g resulting from observing the signal s_g, while $\frac{1}{2}\log\left(\frac{\text{Var}(R_b)}{\text{Var}(R_b|s_b)}\right)$ captures the corresponding reduction in uncertainty about brown return R_b from observing s_b. This constraint binds at equilibrium, so

$$I(R_g, R_b; s_g, s_b) = \kappa.$$

This occurs because the investor optimally uses all available information-processing capacity to maximize expected utility. If the constraint were not binding, the investor could further reduce uncertainty about the returns and improve their portfolio choice, leading to higher utility. Under the Gaussian signal structure introduced in (6), the mutual information (9) admits the closed-form expression (see Appendix A for more details):

$$I(R_g, R_b; s_g, s_b) = \frac{1}{2}\log\left(1 + \frac{\alpha_g^2 \sigma_{R_g}^2}{\sigma_g^2}\right) + \frac{1}{2}\log\left(1 + \frac{\alpha_b^2 \sigma_{R_b}^2}{\sigma_b^2}\right). \tag{10}$$

4.1 Lagrangian Formulation

At time $t = 1$ the investor receives the stochastic income e_1 and chooses first-period consumption c_1. The remaining wealth $e_1 - c_1$ is allocated across the green and brown stocks, and the risk-free asset. At time $t = 2$, the investor receives an additional stochastic amount e_2 and the realized asset returns, and consumes all available wealth. Therefore, the second-period consumption c_2 satisfies

$$c_2 = (e_1 - c_1)R_p + e_2. \tag{11}$$

Note that the wealth at $t = 2$ consists of the realized returns of first-period investment plus second-period income e_2. The investor jointly chooses first-period consumption and portfolio allocation to maximize expected lifetime utility over two periods (2).

Accordingly, the investor solves the following problem:

$$\max_{c_1, \alpha_g, \alpha_b, \alpha_f} \mathbb{E}[u(c_1, E) + \beta u(c_2, E)] \tag{12}$$

subject to:

$$I(R_g, R_b; s_g, s_b) = \kappa \tag{13}$$

$$\alpha_g + \alpha_b + \alpha_f = 1, \tag{14}$$

with $\alpha_g \geq 0,\ \alpha_b \geq 0,\ \alpha_f \geq 0$, and with $0 \leq c_1 \leq e_1$. We have implicitly assumed that $E_1 = E_2 = E$, since the environmental damage per unit of brown stock does not change over time.[4]

[4] Note that environmental impact arises only from asset holdings, which are chosen after consumption at $t = 1$. As a result, first-period utility takes the form $u(c_1, E) = \frac{c_1^{1-\gamma}}{1-\gamma}$, while second-period utility is given by $u(c_2, E) = \frac{c_2^{1-\gamma}}{1-\gamma} - \phi \alpha_b E_b$.

The Lagrangian reads:

$$\mathscr{L} = \mathbb{E}[u(c_1,E) + \beta u(c_2,E)] - \lambda_1 \left(I(R_g,R_b;s_g,s_b) - \kappa\right) - \lambda_2(\alpha_g + \alpha_b + \alpha_f - 1), \tag{15}$$

where $\mathscr{L} := \mathscr{L}(c_1, \alpha_g, \alpha_b, \alpha_f, \lambda_1, \lambda_2)$, and where λ_1, λ_2 are the Lagrange multipliers.

The first order conditions can be written as

$$\begin{cases} c_1^{-\gamma} - \beta \mathbb{E}[c_2^{-\gamma} R_p] = 0 \\ \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_g] - \lambda_1 \frac{\alpha_g \sigma_{R_g}^2}{\sigma_g^2 + \alpha_g^2 \sigma_{R_g}^2} - \lambda_2 = 0, \\ \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_b] - \beta \phi E_b - \lambda_1 \frac{\alpha_b \sigma_{R_b}^2}{\sigma_b^2 + \alpha_b^2 \sigma_{R_b}^2} - \lambda_2 = 0, \\ \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_f] - \lambda_2 = 0, \\ I(R_g,R_b;s_g,s_b) - \kappa = 0, \\ \alpha_g + \alpha_b + \alpha_f - 1 = 0, \end{cases} \tag{16}$$

where

$$I(R_g,R_b;s_g,s_b) = \frac{1}{2}\log\left(1 + \frac{\alpha_g^2 \sigma_{R_g}^2}{\sigma_g^2}\right) + \frac{1}{2}\log\left(1 + \frac{\alpha_b^2 \sigma_{R_b}^2}{\sigma_b^2}\right),$$

and where

$$c_2 = (e_1 - c_1)(\alpha_g R_g + \alpha_b R_b + \alpha_f R_f) + e_2,$$

with $\alpha_g \geq 0,\ \alpha_b \geq 0,\ \alpha_f \geq 0$, and with $0 \leq c_1 \leq e_1$. See Appendix B for more details.

5 Numerical solution

The system of equations in (16) is highly nonlinear and therefore requires numerical methods to be solved. To this end, we specify the model parameters as follows.

5.1 Parameter specification

We adopt the following parameter values:

1. First-period income e_1 is assumed to follow a normal distribution with mean 0.085 and standard deviation 0.023, while second-period income e_2 is normally distributed with mean 0.05 and standard deviation 0.03.

2. The coefficient of relative risk aversion γ in the CRRA utility function (1) is set to 2, a standard and well-established choice in the macro-finance and asset-pricing literature. This value is widely used as a benchmark; for instance, Campbell et al. (2003) adopt $\gamma = 2$ in their strategic asset allocation framework, explicitly characterizing it as representative of a moderate and empirically plausible level of risk aversion.[5]

[5] Mehra and Prescott (1985) argue that values of the coefficient of relative risk aversion, γ, in the range of 2 to 4 are consistent with microeconomic evidence and commonly used macroeconomic calibrations.

3. The parameter ϕ in (1) captures the disutility investors associate with environmental damage, arising specifically from exposure to brown assets. We set this parameter at 0.1 in the model, consistent with Pastor et al. (2021), who choose this value to generate an equilibrium greenium of approximately 1.5% annually. This value is in line with a moderate degree of investor preference for environmental sustainability and is also supported by survey evidence reported in Krueger et al. (2020), which suggests that many investors are willing to accept lower financial returns in exchange for more sustainable investment outcomes. [6]

4. The parameter κ in the information constraint (13) represents the investor's information-processing capacity and serves as a mutual information bound on how much information about asset returns and environmental impacts can be absorbed. We set $\kappa = 0.5$ as a normalized midpoint value, corresponding to a moderate level of attention and cognitive bandwidth. As emphasized by Sims (2003), κ should not be interpreted as a direct empirical measure of cognitive ability, but rather as a modeling device that constrains the rate at which uncertainty can be reduced through information acquisition. [7] Similarly, Van Nieuwerburgh and Veldkamp (2010) normalize the information budget and study equilibrium outcomes under interior levels of information-processing capacity. Furthermore, in the sensitivity analysis presented in subsection 5.3, we explore how variations in the information-processing capacity (between 0 and 1) affect both optimal portfolio allocations and consumption choices.

5. Green and brown asset returns are assumed to follow a normal distribution, namely $R_g \sim \mathcal{N}(\mu_{R_g}, \sigma^2_{R_g})$ and $R_b \sim \mathcal{N}(\mu_{R_b}, \sigma^2_{R_b})$, with $\mu_{R_g} = 1.08$ and $\mu_{R_b} = 1.10$, while the risk-free asset is set as $R_f = 1.02$. These values represent a realistic return hierarchy. Brown stocks offer slightly higher returns to compensate for transition and environmental risks, consistent with the presence of carbon premia (Bolton and Kacperczyk, 2021; Pastor et al., 2021), while green and sustainable assets offer slightly lower financial returns in line with evidence that investors are willing to sacrifice performance for sustainability characteristics (Hartzmark and Sussman, 2019; Pedersen et al., 2021). The risk-free rate of 2% is consistent with long-run real interest rates in developed economies (Del Negro et al., 2019; Rogoff et al., 2024). Moreover, their variances are set at $\sigma^2_{R_g} = 0.15$ and $\sigma^2_{R_b} = 0.2$, and this reflects the riskiness of the investment of the brown stock compared to the green one (Krueger et al., 2020; Bolton and Kacperczyk, 2021).

[6] Specifically, Krueger et al. (2020) show that institutional investors express moderate willingness to sacrifice returns for environmental sustainability, ranking climate risks as financially material but not dominant. This supports setting ϕ at a modest level to reflect realistic ESG/green preferences observed in practice.

[7] Sims (2003) illustrates the mechanism using a range of purely demonstrative capacities, ranging from 0.111 to 3.56 bits per time unit. They are chosen only to show how the dynamics of optimal behavior change under different information-flow regimes. He explicitly notes that even one bit per month is "a very modest requirement" and that actual human information-processing ability is "vastly higher", with effective low capacity arising because individuals typically allocate most of their attention to non-macroeconomic information. Accordingly, following the broader rational-inattention literature, we treat κ as a normalization choice that ensures an interior solution for attention allocation rather than as a parameter with a natural empirical scale or behavioral interpretation. A similar treatment appears in Van Nieuwerburgh and Veldkamp (2010), where the information budget κ is normalized to study how different learning technologies (additive precision vs. entropy-based) shape under-diversification; there too, no attempt is made to estimate the scale of κ, and its magnitude has no empirical content beyond determining the region in which interior vs. corner solutions arise.

6. The signal variances are set at $\sigma_g^2 = 0.04$ and $\sigma_b^2 = 0.06$. These values reflect greater uncertainty in the signals associated with brown stocks, consistent with the idea that brown industries face more legal, regulatory and reputational risk (HSU et al., 2023). Higher noise in brown signals is also in line with empirical findings that ESG-related risks for carbon-intensive firms are difficult to price and hedge due to lack of disclosure and non-standardized data (Krueger et al., 2020). Moreover, Engle et al. (2020) argue that 'climate news contains substantial noise, particularly for carbon-intensive firms'.

7. The parameter E_b takes into account the environmental cost per unit of brown stock and it is set as $E_b = 0.5$. This value is a normalized measure to scale the environmental penalty in utility. It allows the model to capture the trade-off between financial return and environmental damage in a tractable way.

Table 1 summarizes the parameters specification used to numerically solve the non-linear system of equations (16).

Table 1: Parameters' specification for the model.

Parameter	**Specification**
First-period income e_1	$e_1 \sim \mathcal{N}(0.085, 0.023)$
Second-period income e_2	$e_2 \sim \mathcal{N}(0.05, 0.03)$
Coefficient of relative risk aversion γ	$\gamma = 2$
Environmental impact E_b	$E_b = 0.5$
Weight on environmental impact ϕ	$\phi = 0.1$
Information-processing capacity κ	$\kappa = 0.5$
Green returns R_g	$R_g \sim \mathcal{N}(1.08, 0.15)$
Brown returns R_b	$R_b \sim \mathcal{N}(1.10, 0.2)$
Signal noise σ_g^2	$\sigma_g^2 = 0.04$
Signal noise σ_b^2	$\sigma_b^2 = 0.06$

5.2 Solution

To solve the system of nonlinear equations derived from the first-order conditions (16) for c_1 and $\alpha_g, \alpha_b, \alpha_f$, we adopt the Newton-Raphson method. We proceed as follows:

1. Initial Guess. We start with an initial guess for c_1, $\alpha_g, \alpha_b, \alpha_f$, and λ_1, λ_2:

$$c_1^{(0)} = 0.04, \quad \alpha_g^{(0)} = 0.3, \quad \alpha_b^{(0)} = 0.3, \quad \alpha_f^{(0)} = 0.4, \quad \lambda_1^{(0)} = 1, \quad \lambda_2^{(0)} = 0.1.$$

2. Compute Residuals. We evaluate the residuals of the first-order conditions, the information-

processing and budget constraints at the current guess:

$$\begin{aligned}
\text{Residual}_1 &= c_1^{-\gamma} - \beta \mathbb{E}[c_2^{-\gamma} R_p] \\
\text{Residual}_2 &= \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_g] - \lambda_1 \frac{\alpha_g \sigma_{R_g}^2}{\sigma_g^2 + \alpha_g^2 \sigma_{R_g}^2} - \lambda_2 \\
\text{Residual}_3 &= \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_b] - \beta \phi E_b - \lambda_1 \frac{\alpha_b \sigma_{R_b}^2}{\sigma_b^2 + \alpha_b^2 \sigma_{R_b}^2} - \lambda_2 \\
\text{Residual}_4 &= \beta \mathbb{E}[c_2^{-\gamma}(e_1 - c_1) R_f] - \lambda_2 \\
\text{Residual}_5 &= I(R_g, R_b; s_g, s_b) - \kappa \\
\text{Residual}_6 &= \alpha_g + \alpha_b + \alpha_f - 1
\end{aligned}$$

3. Update Guess. We use the Newton-Raphson routine to update the guess:

$$\begin{pmatrix} c_1^{(k+1)} \\ \alpha_g^{(k+1)} \\ \alpha_b^{(k+1)} \\ \alpha_f^{(k+1)} \\ \lambda_1^{(k+1)} \\ \lambda_2^{(k+1)} \end{pmatrix} = \begin{pmatrix} c_1^{(k)} \\ \alpha_g^{(k)} \\ \alpha_b^{(k)} \\ \alpha_f^{(k)} \\ \lambda_1^{(k)} \\ \lambda_2^{(k)} \end{pmatrix} - J^{-1} \begin{pmatrix} \text{Residual}_1 \\ \text{Residual}_2 \\ \text{Residual}_3 \\ \text{Residual}_4 \\ \text{Residual}_5 \\ \text{Residual}_6 \end{pmatrix},$$

where J denotes the Jacobian matrix of the residuals with respect to $c_1, \alpha_g, \alpha_b, \alpha_f$ and λ_1, λ_2.

4. Check Convergence. We repeat the process until the residuals are smaller than a predefined tolerance (e.g. $\text{tol} = 10^{-6}$).

Assuming the parameters' specification given in Table 1, the algorithm converges robustly to the solution

$$c_1^* \approx 0.067, \quad \alpha_g^* \approx 0.29, \quad \alpha_b^* \approx 0.57, \quad \alpha_f^* \approx 0.14. \tag{17}$$

5.3 Sensitivity analysis

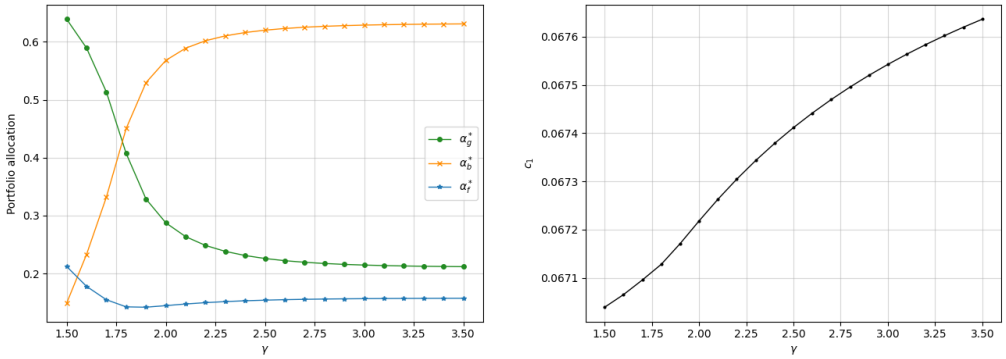
We perform a sensitivity analysis of the more impactful parameters of the system described by (16), to measure their effect on the optimal allocations $\alpha_g^*, \alpha_b^*, \alpha_f^*$ and on the optimal consumption c_1^*. We start with a complete one-parameter sensitivity analysis to check the overall impact of each single parameter, and then perform a two-parameter sensitivity analysis to gain deeper insights.

5.3.1 One-parameter sensitivity analysis of optimal portfolio weights

To assess the overall impact of each parameter, we construct an evenly spaced grid across its plausible range and vary that parameter over the grid while holding all other parameters fixed, as specified in Table 1. This approach allows us to isolate the marginal effect of each parameter on optimal choices.

Figure 1 illustrates how varying the relative risk aversion parameter γ in (16) affects the optimal allocations $\alpha_g^*, \alpha_b^*, \alpha_f^*$ and initial consumption c_1^*. To do so, we use an equispaced grid from $\gamma = 1.5$ to $\gamma = 3.5$, with a step of 0.1. As γ increases, investors allocate more to brown assets and less to green assets. The allocation to the risk-free asset remains largely stable, exhibiting a slight decline as γ rises from 1.50 to 1.75 (blue dotted line). The right panel of Figure 1 shows that optimal initial consumption rises gradually with increasing risk aversion. This pattern reflects a preference among more risk-averse individuals for immediate, certain consumption over uncertain future returns. The measured, steady increase suggests that while higher risk aversion encourages greater present consumption, its impact on consumption decisions remains gradual, highlighting a consistent trade-off between current utility and future wealth accumulation.

Figure 1: Sensitivity of Portfolio Allocations to Changes in γ

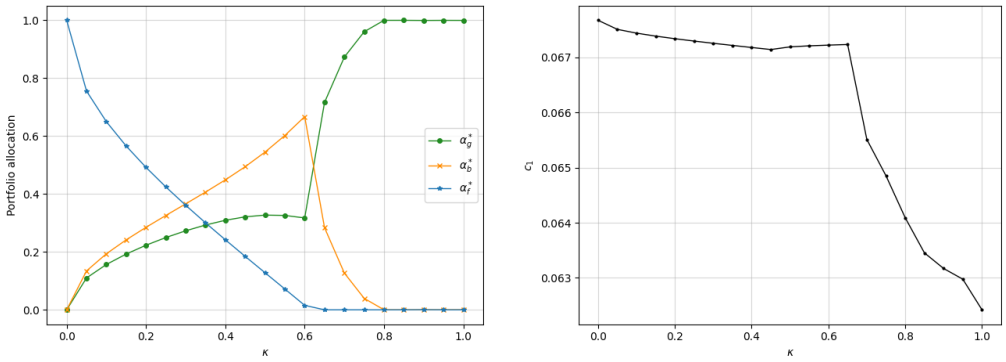


Notes: The left panel shows how the three asset-class allocations vary with γ, while the right panel illustrates the effect of γ on initial consumption c_1.

Figure 2 illustrates how variations in the information-processing capacity parameter, κ, in (16) influence both the optimal portfolio allocations (left panel) and the optimal initial consumption choice (right panel). The parameter κ is assumed to vary from 0 to 1 in increments of 0.05. At the lower bound, $\kappa = 0$, investors are unable to extract any information from the available signals. Lacking informative content about risky returns, they optimally allocate their entire wealth to the risk-free asset, fully avoiding exposure to risk. As κ increases, investors gain a greater ability to process and extract useful information from the available signals. This enhanced informational capacity allows them to form more precise expectations about asset returns and risks, which in turn leads to systematic adjustments in both portfolio composition and consumption decisions. In particular, higher values of κ translate into a progressively larger allocation to green stocks, accompanied by declining positions in brown stocks and the risk-free asset. This reallocation reflects the fact that, with more information at their disposal, investors place greater confidence in the superior expected returns and environmental attributes of green assets and, therefore, are more willing to tilt their portfolios in that direction.

The effect becomes especially pronounced as κ approaches one, corresponding to the case of full information-processing capacity. In this limit, the investor fully internalizes the informa-

Figure 2: Sensitivity of Portfolio Allocations to Changes in κ



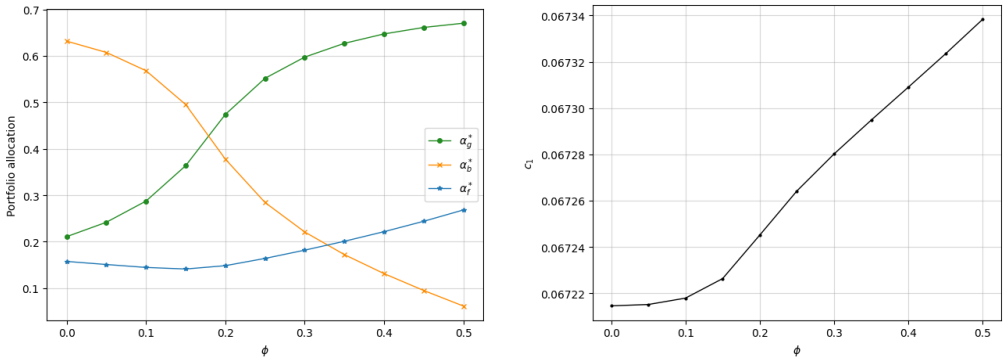
Notes: The left panel shows how the three asset-class allocations vary with κ, while the right panel illustrates the effect of κ on initial consumption c_1.

tional content of the signals and thus allocates almost all wealth to green assets, which become overwhelmingly preferred. Conversely, the shares invested in the brown asset and the risk-free asset shrink virtually toward zero. Optimal consumption similarly declines, as investors choose to allocate nearly all available resources to the green asset in anticipation of higher future payoffs. This pattern highlights how improvements in information-processing ability can fundamentally reshape investment behavior, amplifying both environmental and financial motivations for green-oriented portfolio choices.

Figure 3 shows how changes in the weight assigned to adverse environmental impact, ϕ, in (16) influence the optimal portfolio allocations and the initial consumption choice. We vary ϕ from 0 to 0.5 in increments of 0.05. As ϕ increases, investors place progressively greater emphasis on environmental harm when making decisions. This heightened concern leads to a clear rebalancing of the portfolio: the allocation to brown assets steadily declines, while the allocation to green assets rises. In other words, stronger environmental preferences push investors to shift their wealth away from environmentally damaging assets and toward greener alternatives. The allocation to the risk-free asset exhibits a mild increase at higher values of ϕ, suggesting a small precautionary component as environmental preferences intensify. The optimal initial consumption also rises with ϕ (Figure 3, right panel). This indicates that when investors care more about environmental impact, they choose to consume more today while still reallocating their remaining wealth toward greener assets. Overall, increasing ϕ amplifies both the preference for environmentally friendly investments and a tendency toward higher current consumption.

Figures 4 and 5 illustrate how variations in signal noise influence optimal allocations and initial consumption. Figure 4, in particular, explores the consequences of increasing noise in the green signal. A higher value of σ_g^2 – that is, a noisier green signal – reduces the investor's ability to infer returns information from it. Because attention is limited and must be directed where it yields the greatest benefit, investors reallocate attention away from the green signal. Consequently, the optimal allocation to green stocks decreases, while the allocation to brown stocks increases. Notably, once the noise exceeds a certain threshold ($\sigma_g^2 > 0.1$), the share of

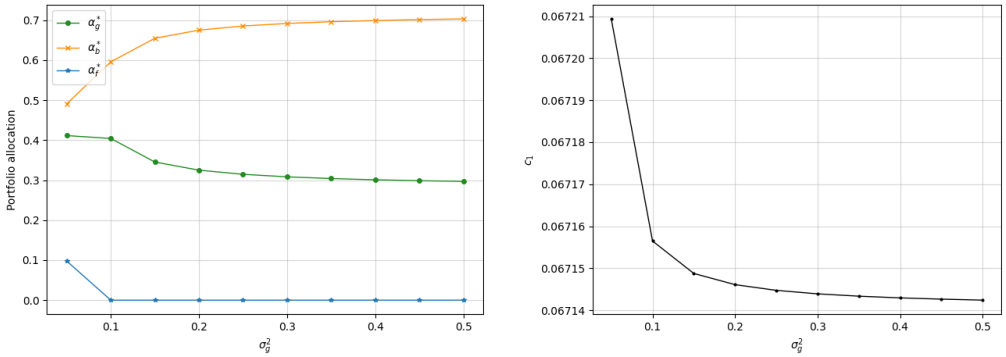
Figure 3: Sensitivity of Portfolio Allocations to Changes in ϕ



Notes: The left panel shows how the three asset-class allocations vary with ϕ, while the right panel illustrates the effect of ϕ on initial consumption c_1.

wealth allocated to the green asset drops by roughly 10 percentage points (Figure 4, green line). Initial optimal consumption also declines, which indicates that investors choose to save more to buffer against the heightened uncertainty in the risky portfolio.

Figure 4: Sensitivity of Portfolio Allocations to Changes in σ_g^2



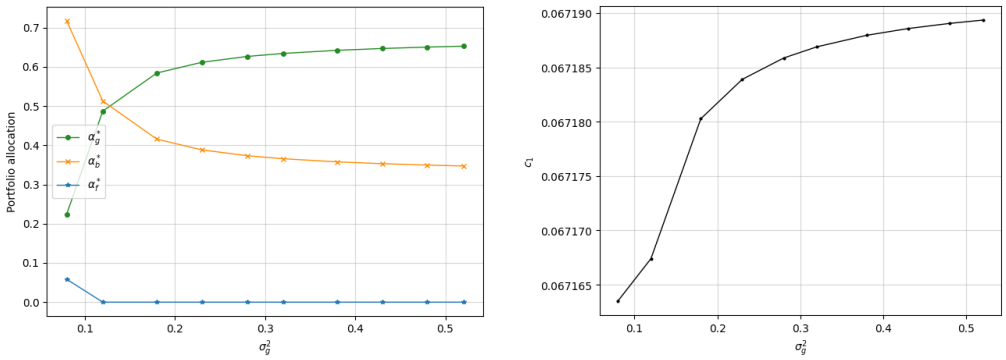
Notes: The left panel shows how the three asset-class allocations vary with σ_g^2, while the right panel illustrates the effect of σ_g^2 on initial consumption c_1.

Similarly, Figure 5 shows that a higher σ_b^2 reduces the investor's ability to extract returns information from the brown signal, which makes brown stocks relatively less attractive than green stocks. As a result, the optimal allocation to brown assets declines, while the allocation to green assets increases. In this case, initial optimal consumption rises (Figure 5, right panel), reflecting the fact that, with fewer appealing risky investment opportunities, investors choose to consume more in the present rather than saving for a less rewarding portfolio.

Table 2 summarizes the effects of a change in the parameters in the allocations of the three asset and on the initial consumption c_1.

The quantitative results highlight a sharp asymmetry in green versus brown investing that arises from the interplay of environmental preferences, risk aversion, and, most crucially, limited

Figure 5: Sensitivity of Portfolio Allocations to Changes in σ_b^2



Notes: The left panel shows how the three asset-class allocations vary with σ_b^2, while the right panel illustrates the effect of σ_b^2 on initial consumption c_1.

Table 2: Sensitivity Analysis Summary for Initial Consumption and Portfolio Allocation

Parameter	**Effect on α_g^***	**Effect on α_b^***	**Effect on α_f^***	**Effect on c_1^***
$\gamma \uparrow$	Decreases	Increases	Negligible	Increases
$\kappa \uparrow$	Increases	Decreases (eventually)	Decreases	Decreases
$\phi \uparrow$	Increases	Decreases	Increases	Increases
$\sigma_g^2 \uparrow$	Decreases	Increases	Decreases	Decreases
$\sigma_b^2 \uparrow$	Increases	Decreases	Decreases	Increases

information-processing capacity. Higher risk aversion shifts portfolios toward brown assets because, under limited attention, investors prioritize assets with higher expected returns that compensate for residual uncertainty. At the same time, the risk-free allocation remains broadly stable. Stronger environmental preferences, by contrast, directly increase demand for green and risk-free assets. They also crowd out brown assets, even when brown stocks initially offer higher average returns.

The central driver of sustainable portfolio reallocation, however, is the investor's information-processing capacity κ. At low levels of κ, severe rational inattention keeps the posterior uncertainty high for both green and brown assets, and green allocations remain modest: investors simply cannot discern the lower long-term environmental and regulatory risk embedded in green firms or the hidden transition and stranded-asset risks of brown firms. As κ rises, the investor endogenously devotes a disproportionate share of the newly available information-processing capacity to the green asset. This occurs because, as information accumulates, learning about brown assets increasingly reveals transition and regulatory risks that reduce their expected utility once environmental preferences are taken into account. Consequently, the marginal value of further attention to brown assets declines, and additional capacity is optimally redirected toward green returns. As posterior uncertainty falls, green assets are perceived as relatively less risky in utility terms, while brown assets become less attractive once both financial and environmen-

tal risks are fully internalized. The resulting re-allocation is highly non-linear; green portfolio weights rise sharply beyond a critical κ-threshold, whereas brown holdings collapse (often approaching zero), and the demand for the risk-free asset declines as confidence in informed risky positions increases. Worse signal quality reinforces these patterns asymmetrically: higher noise in green signals directly depresses green investment, whereas higher noise in brown signals often increases green holdings by leaving investors precautionary about undisclosed climate risks in the brown sector.

In short, limited attention could be seen as the key bottleneck slowing the private-sector green transition. Expanding investors' effective processing capacity κ, through better disclosure, standardized ESG data, or financial education, triggers a rapid and powerful shift of capital from brown to green assets, even among moderately environmentally conscious investors. Rational inattention thus transforms environmental preferences from a mild tilt into a decisive force for portfolio decarbonization once cognitive constraints are relaxed.

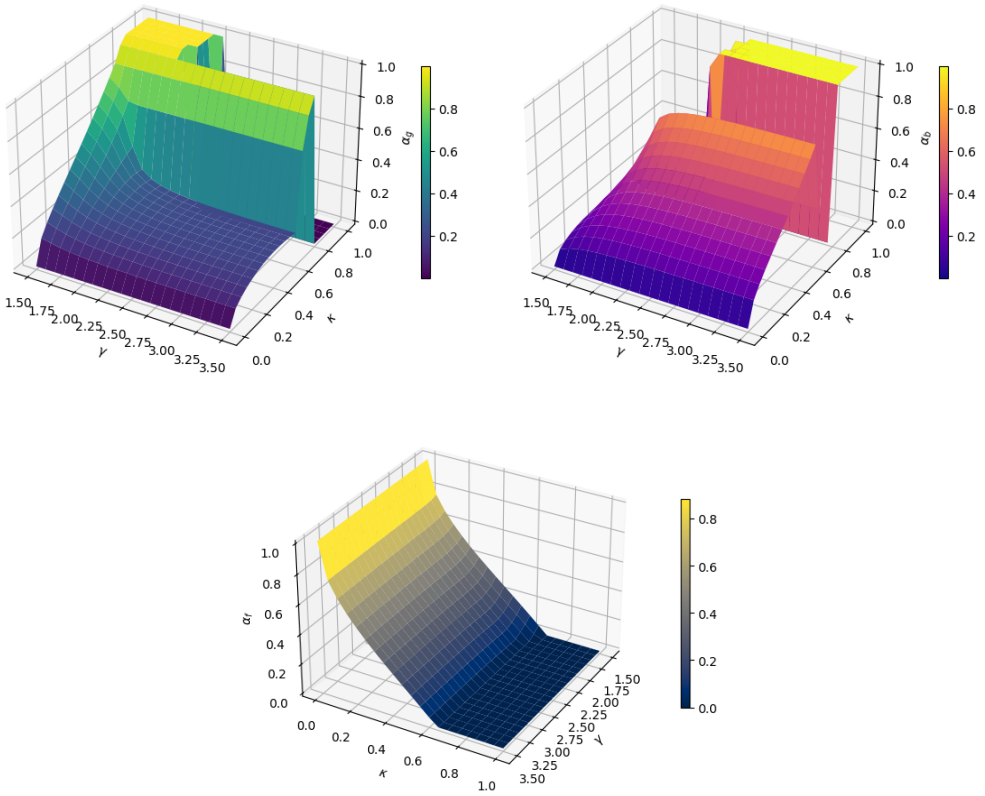
5.3.2 Two-parameter sensitivity analysis of optimal portfolio weights

To gain deeper insights into the mechanisms of the model, we conduct a series of two-dimensional sensitivity analyses. Specifically, we examine how optimal portfolio weights α_g^*, α_b^*, and α_f^* respond to joint variation in information-processing capacity κ and three key drivers: (i) risk aversion γ, (ii) green-signal noise variance σ_g^2, and (iii) brown-signal noise variance σ_b^2. For each parameter pair $(\kappa, \cdot)$, we construct a fine, evenly spaced grid over economically plausible ranges (as reported in Table 1). While one parameter varies along its grid, the other is held fixed at its baseline value. This approach allows us to visualize interaction effects and non-linearities that are masked in one-dimensional exercises. The resulting heatmaps and contour plots reveal sharp threshold behavior and strong asymmetries, particularly the pronounced non-linear surge in green holdings and collapse of brown positions once κ crosses critical levels, confirming that limited attention is the primary bottleneck governing the speed and extent of portfolio decarbonization.

Figure 6 shows the joint impact of information-processing capacity κ and risk aversion γ on optimal portfolio weights. When risk aversion is low, the green weight α_g^* rises steadily and monotonically with κ, eventually saturating near its feasible upper bound. Ample attention allows investors to sharply reduce posterior uncertainty about green returns and to fully exploit the lower fundamental risk that becomes visible once climate-related tail risks are properly measured, leading to aggressive green positions. As risk aversion increases (from $\gamma \approx 1.75$), the relationship turns markedly non-monotonic and reveals a distinctive hump-shaped pattern in κ. At moderate-to-high γ, the green weight initially increases only modestly with κ because the investor remains cautious about the (ex ante) higher volatility of green assets. Beyond a critical threshold of attention (i.e., around $\kappa > 0.8$), however, posterior variance collapses once sufficient information is processed, triggering a sudden and steep surge in α_g^*. However, at high levels of information-processing capacity, posterior uncertainty about risky returns becomes negligible and the risk-free asset is endogenously crowded out. Portfolio choice then converges to a risky corner solution. Which risky asset dominates depends on the balance between ex-

pected returns and non-pecuniary environmental costs. For sufficiently high risk aversion, the green asset no longer enjoys a perceived risk advantage once information is abundant, and the higher expected return on the brown asset outweighs the linear environmental penalty. As a result, the optimal portfolio switches discontinuously from green to brown, producing the sharp collapse in green holdings observed in Figure 6.

Figure 6: Optimal Portfolio Weights across Varying Information Capacity κ and Risk Aversion γ.



Notes: Each panel shows one allocation: top-left α_g, top-right α_b, and bottom α_f. Color bars give the magnitude of the allocation. All other parameters are set to their baseline values.

The brown asset allocation α_b^* exhibits the mirror-image pattern of the green weight. It remains relatively high when information capacity κ is low (investors cannot discern brown climate risks and therefore underestimate their true exposure) and, conditional on sufficiently high information-processing capacity, when risk aversion γ is extremely high. In this latter case, posterior uncertainty is largely eliminated, the risk-free asset is crowded out, and portfolio choice converges to a risky corner solution in which the higher expected return on brown assets outweighs the linear environmental penalty. Conversely, α_b^* collapses – often abruptly to near zero – precisely where the green weight surges: once κ crosses the threshold at which investors can accurately reveal and price the large transition, regulatory, and stranded-asset risks embedded in brown firms. These knife-edge transitions coincide with the points where interior solutions give way to binding corner solutions ($\alpha_b^* = 0$ or α_g^* close to its feasible upper bound), producing the

starkly contrasting heatmaps in Figure 6, left and right panels. Thus, the same attentional phase transition that unlocks aggressive green investment simultaneously drives the decisive exclusion of brown assets from the optimal portfolio.

Our sensitivity analysis reveals a sharp, non-linear reallocation toward green assets once investors' information-processing capacity exceeds a critical threshold. This tipping-point mechanism offers a microfoundation for the sudden green outperformance documented by Ardia et al. (2023) upon the materialization of climate shocks. While Ardia et al. (2023) attribute these return patterns to sudden spikes in climate concerns, our model demonstrates that such shifts emerge endogenously from relaxed cognitive constraints, rather than purely behavioral reactions. Specifically, limited attention hinders the timely integration of ESG/green signals into portfolio choices; however, as capacity increases, investors swiftly reprice climate risks, prompting rapid divestment from brown assets and a pronounced overweighting of green ones.

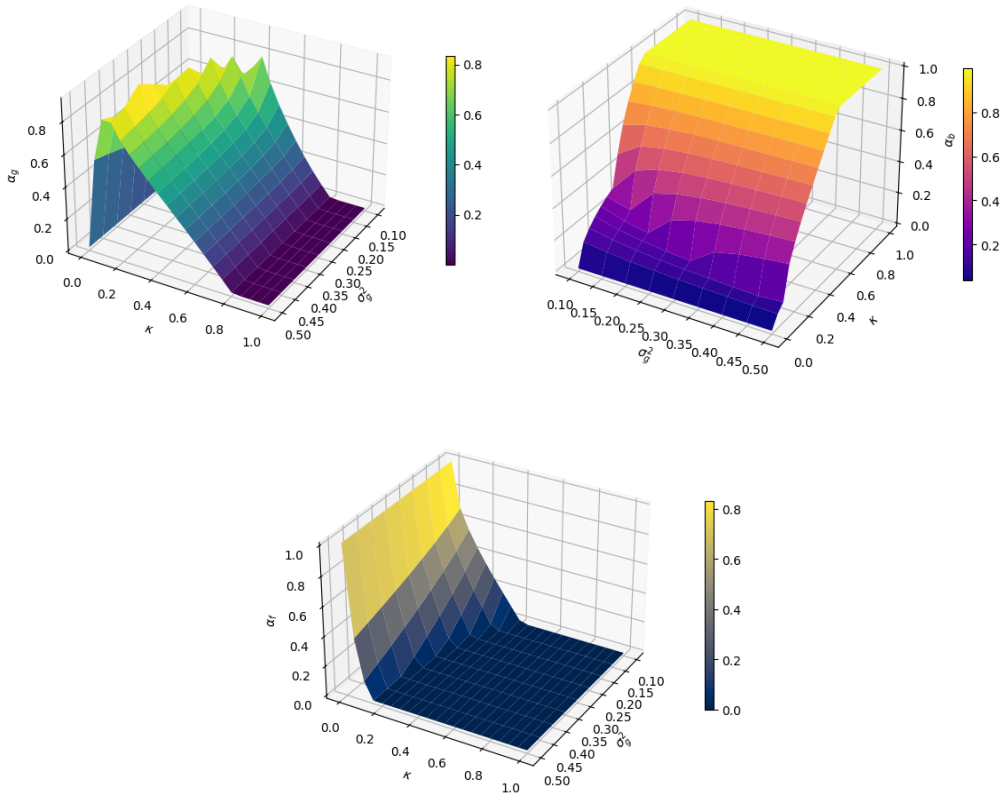
By contrast, the risk-free allocation α_f^* displays a smooth and nearly monotonic decline toward zero as information-processing capacity κ increases, and this pattern holds across almost the entire range of risk aversion γ. Higher κ enables investors to substantially reduce posterior uncertainty about both green and brown returns, thereby lowering the perceived risk of the risky portfolio as a whole. Even highly risk-averse investors, who would otherwise remain parked in the risk-free asset, find it optimal to take informed positions in at least one of the risky assets once sufficient precision is achieved. The gradual substitution out of the risk-free asset therefore reflects a pure confidence effect: greater effective capacity translates directly into greater willingness to bear risk, with the composition of that risk (green versus brown) determined by the mechanisms described earlier, but the overall exposure to risk rising steadily with κ (see Figure 6, bottom panel).

Figure 7 illustrates the joint effects of information-processing capacity κ and green-signal noise variance σ_g^2 (where higher σ_g^2 means lower exogenous signal precision) on optimal portfolio weights. As κ rises, the risk-free weight α_f^* declines smoothly and rapidly toward zero at all levels of σ_g^2 (bottom panel). Greater capacity allows investors to acquire precise posterior beliefs regardless of the raw signal quality, thereby reducing overall portfolio risk and encouraging substitution out of the safe asset into risky assets. The speed of this exodus from the risk-free asset, however, depends critically on green-signal noise. When the green signal is very noisy, extracting useful information is costly in entropy terms. However, because attention must be allocated to financial-decision making, marginal increases in κ are disproportionately spent attempting to reduce green uncertainty, even though precision gains are gradual.

Posterior uncertainty about green returns falls once κ is sufficiently high, perceived risk drops sharply, and investors aggressively shift into green stocks. Consequently, α_f^* drops almost immediately; (ii) when the green signal is already precise (low σ_g^2), additional capacity yields diminishing marginal reductions in posterior variance. The confidence gain is more gradual, and the decline of α_f^* is correspondingly slower and smoother. Thus, counterintuitively, poorer exogenous green-signal quality accelerates the abandonment of the risk-free asset—because scarce attention is forced to work harder on the green dimension, delivering larger early gains in perceived Sharpe ratio and triggering a faster tilt toward informed risky positions.

A striking asymmetry emerges in which risky asset absorbs the portfolio share released from

Figure 7: Optimal Portfolio Weights across Varying Information Capacity κ and of the Green Signal σ_g^2



Notes: Each panel shows one allocation: top-left α_g, top-right α_b, and bottom α_f. Color bars give the magnitude of the allocation. All other parameters are set to their baseline values.

the risk-free asset, and this choice hinges crucially on the exogenous quality of the green signal (Figure 7, top-left and top-right panels). When the green signal is already precise (low σ_g^2), the investor can achieve high posterior precision about the green asset even with very modest attention. At low κ, the green weight α_g^* therefore rises early and substantially (top-left corner of the top panel). As κ increases further, however, marginal attention becomes relatively scarce for the green dimension (little precision left to gain), so the investor reassigns newly available capacity to the brown asset. Once brown climate risks are revealed, the brown weight surges while the green position is crowded out entirely, producing the dramatic vertical cliff in α_g^* at intermediate-to-high κ. When the green signal is noisy (high σ_g^2), useful information about the green asset is expensive to extract. Even moderate attention yields only limited precision gains, so the green asset remains perceived as risky and unattractive at virtually all levels of κ. Investors instead direct almost all capacity toward the brown signal (which is assumed cleaner in the calibration), quickly uncover brown environmental risks, and abandon the brown asset early. The portfolio then settles into a combination of risk-free and whatever small residual green exposure is tolerated, with α_b^* collapsing early and α_g^* staying near zero throughout. Thus, paradoxically, a naturally precise green signal can backfire at high capacity: it causes initial

overinvestment in green, followed by a sharp reversal once attention is redirected to exposing brown risks. A persistently noisy green signal, by contrast, prevents any meaningful green allocation from emerging in the first place, locking the portfolio into brown exclusion and risk-free dominance. These results underscore that transparency and verifiability of green claims (low σ_g^2) are essential for sustained – not just transitory – green investment.

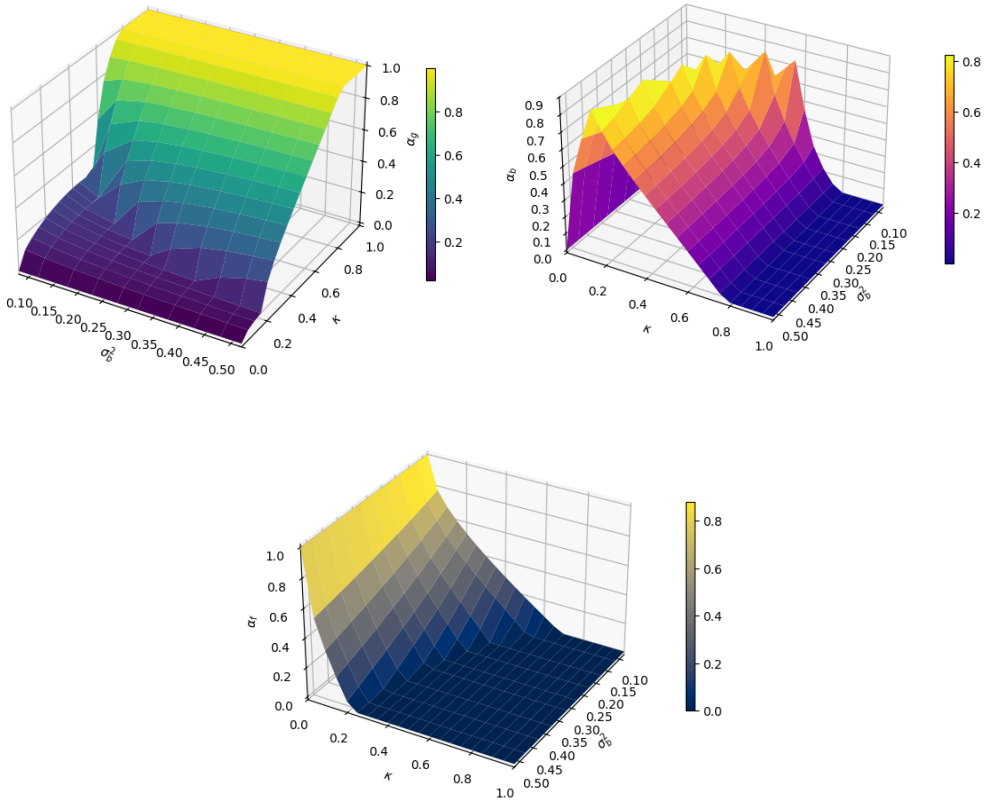
Figure 8 displays the joint variation of optimal allocations with information capacity κ and brown-signal noise variance σ_b^2 (higher σ_b^2 implies lower exogenous precision). The patterns are the exact mirror image of Figure 7, with green and brown roles reversed. When the brown signal is clean (low σ_b^2), the brown asset absorbs most of the portfolio share freed from the risk-free asset even at modest levels of κ. Posterior precision about brown returns is achieved cheaply, so α_b^* rises early and remains high across most of the κ-space (see top-right panel). Only at very high κ ($\kappa > 0.6$) does the investor begin redirecting marginal attention toward the green asset, eventually eroding the brown position. Conversely, when the brown signal is noisy (high σ_b^2), extracting information about brown returns becomes costly. The investor therefore devotes the bulk of any available capacity to the brown dimension in order to reveal its large hidden climate and transition risks. As a result, α_b^* collapses to zero extremely rapidly – often at quite low values of κ – producing a steep cliff at the left-hand side of the surface (top-right panel). The green asset then captures the remaining risky allocation, and the portfolio tilts decisively toward green. The risk-free weight α_f^* again declines monotonically with κ, but the exit from the safe asset is noticeably slower and more gradual when the brown signal is precise. This reflects the environmental taste penalty (i.e., $-\phi E_b$) attached to brown returns: even with perfect information, the brown asset is less attractive than a comparable green asset, so investors require higher overall capacity before they are willing to abandon the risk-free asset entirely. In short, a greater opacity of brown risks (high σ_b^2) acts as a powerful accelerator of portfolio decarbonization: it forces early and aggressive attention to the brown sector, rapidly exposes its unsustainable status and triggers a swift exodus from brown assets to green ones. Clean brown signals, by contrast, delay the green transition by keeping brown assets artificially viable for longer. This asymmetry demonstrates that mandatory climate-risk disclosure targeted specifically at brown firms (i.e., policies that reduce σ_b^2) would dramatically hasten capital reallocation toward sustainability.

5.4 Additional analyses

We now sharpen the analysis with two targeted extensions that test whether the green-vs-brown divide becomes even more dramatic.

In the first extension, we fix intermediate consumption c_1 at a constant level. This modification may lead to a significantly sharper and potentially more extreme separation between green and brown holdings than in the baseline. By removing the consumption-saving margin and the associated wealth effects, every marginal increase in information capacity κ is fully channeled into refining beliefs about the two risky assets. As a result, attention thresholds are crossed earlier, perceived risk differentials widen more dramatically, and the portfolio may shift abruptly toward corner solutions – either near-complete divestment from brown assets or a strong overweight in green – often at lower levels of κ than under flexible consumption.

Figure 8: Optimal Portfolio Weights across Varying Information Capacity κ and of the Green Signal σ_b^2



Notes: Each panel shows one allocation: top-left α_g^*, top-right α_b^*, and bottom α_f^*. Color bars give the magnitude of the allocation. All other parameters are set to their baseline values.

In the second extension, we replace the linear environmental taste term with a convex quadratic disutility $\phi(\alpha_b E_b)^2$. This change may trigger far more aggressive brown exclusion than the linear specification would suggest. Because the marginal moral, social, reputational, and regulatory cost now rises with the square of the brown weight, even modest confirmed environmental damage E_b can make any positive brown position intolerably expensive. Consequently, as soon as information capacity κ becomes sufficient to credibly reveal the extent of brown externalities, the optimal brown weight may collapse more rapidly to zero, while the green asset captures nearly the entire risky allocation. Together, these two extensions may transform what appears in the baseline as a gradual, cautious green tilt into a swift and decisive decarbonization of the portfolio the moment the attention bottleneck is relaxed.

Both extensions strengthen the central message of the paper: limited attention is the true binding constraint on portfolio decarbonization. Relax that constraint even slightly, and the interplay of precise information with realistic (convex) ethical costs converts a modest, hesitant green tilt into a swift and near-complete exodus from brown assets — a decisive, large-scale reallocation of capital toward sustainability. The green transition is therefore not a slow cultural

evolution; it is a latent phase transition awaiting the release of the cognitive bottleneck.

Fixed consumption. The derivation of the model, the characterization of the optimal allocation, and a series of sensitivity analyses under this alternative specification are presented in Appendix C.

In the spirit of the sensitivity analyses presented earlier, Figure C.1 reports the response of optimal allocations α_g^*, α_b^*, and α_f^* to the main parameters when intermediate consumption c_1 is fixed at $0.7e_1$ (a value close to the endogenous optimum in the baseline calibration). The overall pattern is remarkably similar to the flexible-consumption case (Figures 1–5), confirming the robustness of the core mechanisms. The only notable difference arises in the dynamics with respect to information capacity κ (Figure C.1, panel b): brown holdings α_b^* now begin their abrupt collapse at substantially lower values of κ (around $\kappa \approx 0.45$ instead of the higher 0.6 threshold observed under endogenous consumption). This sharper and earlier divestment from the brown asset is exactly what the theory predicts once the saving margin is removed: without the possibility of adjusting c_1 to smooth wealth fluctuations, every marginal nat of attention is immediately translated into more precise pricing of climate risks, accelerating the revelation of brown externalities and triggering a faster flight to green and risk-free assets. Thus, even a modest restriction on consumption flexibility is sufficient to bring forward the portfolio-decarbonization tipping point by a non-trivial amount.

Non-linear environmental damage. The complete derivation of the model with quadratic brown-investment disutility, together with the closed-form expressions for the optimal allocations, is presented in Appendix D. Figures D.1 through D.5 report the corresponding sensitivity analyses.

Figure D.1 shows the response of optimal allocations α_g^*, α_b^*, α_f^* and consumption c_1^* to varying risk aversion γ. The qualitative pattern remains consistent with the linear-penalty case (Figure 1), but the separation between green and brown assets is far more dramatic, especially at low-to-moderate levels of γ, where the convex moral cost rapidly drives brown holdings to zero. Figure D.2 traces the effect of information-processing capacity κ. The trajectory closely tracks the linear benchmark (Figure 2), yet the tipping point arrives noticeably earlier: the brown weight collapses and the green weight surges at lower values of κ, confirming that quadratic disutility amplifies the decarbonization trigger. Figure D.3 examines the strength of the environmental preference parameter ϕ. Here the departure from the linear case (Figure 3) is most visible: instead of the sharp, near-discontinuous jumps observed under linear taste, both green and brown allocations adjust more gradually and smoothly. The convex penalty spreads the moral cost across allocation levels, softening extreme corner solutions and producing richer interior dynamics while still ultimately favoring green dominance at higher ϕ. Finally, Figures D.4 and D.5 display the impact of signal noise σ_g^2 and σ_b^2. The responses of allocations and consumption remain qualitatively and quantitatively very close to their linear-penalty counterparts (Figures 4–5), underscoring that the core attention-allocation mechanism is robust to the functional form of the environmental cost.

Taken together, the quadratic specification preserves all key insights of the baseline model

while revealing that realistic convexity in ethical costs can either accelerate abrupt decarbonization (when attention or preferences are the binding margins) or generate smoother, more gradual transitions (when the moral penalty itself is the primary driver).

6 Conclusion

This paper develops a rationally inattentive portfolio-choice model to explain how cognitively constrained investors allocate capital between green and brown assets. Building on the mutual-information framework of Sims (2003) and extending the two-period consumption-portfolio setting in the spirit of Batchuluun et al. (2019), we introduce two key features: (i) a finite information-processing capacity κ that limits the investor's ability to learn about returns and environmental externalities, and (ii) direct taste (or distaste) for the sustainability characteristics of the portfolio.

The model is motivated by robust empirical evidence that investors value sustainability (see, among others, Krueger et al., 2020; Pastor et al., 2021), yet systematically underweight green assets and respond sluggishly to climate-policy shocks – patterns that standard rational-expectations models struggle to reconcile. By endogenizing attention allocation via a mutual-information constraint and allowing noisy signals about both financial payoffs and environmental impact, our framework delivers sharp, testable predictions.

Limited processing capacity κ emerges as the central friction. It generates slow incorporation of green-policy news into asset prices, large cross-sectional heterogeneity in portfolio greening, and precautionary underinvestment in green assets despite strong stated sustainability preferences. We show that decisive brown divestment and green overweight occur only when κ exceeds critical thresholds, and that these thresholds decline sharply with stronger environmental preferences ϕ, higher signal precision, or convex moral costs on brown exposure.

The model rationalizes several stylized facts in sustainable finance: persistent underdiversification into green assets, delayed price reactions to climate-related shocks, and the widespread disconnect between investors' expressed ESG concerns and their actual holdings. More broadly, it offers a tractable laboratory for evaluating the effectiveness of disclosure mandates, carbon pricing, and green subsidies – policies that primarily operate by easing the attention bottleneck or by making brown risks costlier to ignore.

By integrating rational inattention with environmental preferences, this paper bridges the sustainable-finance literature and the growing application of information-choice theory to asset pricing, delivering new micro-foundations for the still-sluggish transition toward a greener financial system.

Promising directions for future research include generalizing the preference specification (e.g., to recursive or quadratic utility) and relaxing the independence assumption between green and brown returns. These extensions would shed light on how risk aversion, intertemporal substitution, and information frictions jointly shape portfolio decisions. In particular, allowing for correlated returns would generate a novel interplay between correlation risk, learning, and diversification incentives, with potentially first-order effects on optimal holdings. We defer a comprehensive treatment of these richer environments to future work.

Bibliography

Ardia, D., Bluteau, K., Boudt, K., Inghelbrecht, K., 2023. Climate change concerns and the performance of green vs. brown stocks. Management Science 69, 7607–7632.

Avramov, D., Lioui, A., Liu, Y., Tarelli, A., 2025. Dynamic esg equilibrium. Management Science 71, 2867–2889.

Batchuluun, A., Luo, Y., Young, E. R., 2019. Portfolio choice with information-processing limits. Annals of Economics and Finance 20, 137–162.

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

Campbell, J. Y., Chan, Y. L., Viceira, L. M., 2003. A multivariate model of strategic asset allocation. Journal of Financial Economics 67, 41–80.

Del Negro, M., Giannone, D., Giannoni, M. P., Tambalotti, A., 2019. Global trends in interest rates. Journal of International Economics 118, 248–262.

Dimson, E., KarakaÅŸ, O., Li, X., 2015. Active ownership. The Review of Financial Studies 28, 3225–3268.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., Stroebel, J., 2020. Hedging climate change news. The Review of Financial Studies 33, 1184–1216.

Giglio, S., Kelly, B., Stroebel, J., 2021. Climate finance. Annual Review of Financial Economics 13, 15–36.

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

Hartzmark, S. M., Sussman, A. B., 2019. Do investors value sustainability? a natural experiment examining ranking and fund flows. The Journal of Finance 74, 2789–2837.

Hirshleifer, D., Lim, S. S., Teoh, S. H., 2011. Limited investor attention and stock market misreactions to accounting information. The Review of Asset Pricing Studies 1, 35–73.

HSU, P.-H., LI, K., TSOU, C.-Y., 2023. The pollution premium. The Journal of Finance 78, 1343–1392.

Kacperczyk, M., Van Nieuwerburgh, S., Veldkamp, L., 2016. A rational theory of mutual funds' attention allocation. Econometrica 84, 571–626.

Krueger, P., Sautner, Z., Starks, L. T., 2020. The importance of climate risks for institutional investors. The Review of Financial Studies 33, 1067–1111.

Mackowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. American Economic Review 99, 769–803.

Mehra, R., Prescott, E. C., 1985. The equity premium: A puzzle. Journal of Monetary Economics 15, 145–161.

Pastor, L., Stambaugh, R. F., Taylor, L. A., 2021. Sustainable investing in equilibrium. Journal of Financial Economics 142, 550–571.

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

Peng, L., Xiong, W., 2006. Investor attention, overconfidence and category learning. Journal of Financial Economics 80, 563–602.

Rogoff, K. S., Rossi, B., Schmelzing, P., 2024. Long-run trends in long-maturity real rates, 1311–2022. American Economic Review 114, 2271–2307.

Sims, C. A., 2003. Implications of rational inattention. Journal of Monetary Economics 50, 665–690.

van der Ploeg, F., 2025. Why green subsidies are preferred to carbon taxes: Climate policy with heightened carbon tax salience. Journal of Environmental Economics and Management 130, 103129.

Van Nieuwerburgh, S., Veldkamp, L., 2010. Information acquisition and under-diversification. The Review of Economic Studies 77, 779–805.

A Derivation of the rational inattention's formula in the Gaussian case

Proposition A.1. *Let the asset returns and signals satisfy*

$$R_g \sim \mathcal{N}(\mu_g, \sigma^2_{R_g}), \quad R_b \sim \mathcal{N}(\mu_b, \sigma^2_{R_b}),$$

and

$$s_g = \alpha_g R_g + \varepsilon_g, \quad s_b = \alpha_b R_b + \varepsilon_b,$$

where $\varepsilon_g \sim \mathcal{N}(0, \sigma^2_g)$ *and* $\varepsilon_b \sim \mathcal{N}(0, \sigma^2_b)$*, with* (R_g, ε_g) *independent of* (R_b, ε_b)*. Then the mutual information* $I(R_g, R_b; s_g, s_b)$ *can be expressed as in equation* (10)*.*

Proof. From equation (8), the mutual information is defined as

$$I(R_g, R_b; s_g, s_b) = H(R_g, R_b) - H(R_g, R_b | s_g, s_b),$$

where $H(R_g, R_b)$ is the joint entropy of the returns, and $H(R_g, R_b | s_g, s_b)$ is the conditional entropy of the returns given the signals. Under the independence assumptions, the joint entropy decomposes additively

$$H(R_g, R_b) = H(R_g) + H(R_b), \quad H(R_g, R_b | s_g, s_b) = H(R_g | s_g) + H(R_b | s_b).$$

Hence

$$I(R_g, R_b; s_g, s_b) = H(R_g) - H(R_g | s_g) + H(R_b) - H(R_b | s_b). \tag{A.1}$$

It is known that if (X, Y) are jointly Gaussian, then $H(X) = \frac{1}{2}\log(2\pi e \mathrm{Var}(X))$ and $H(X|Y) = \frac{1}{2}\log(2\pi e \mathrm{Var}(X|Y))$. Hence, we can apply it to equation (A.1), and this leads to

$$\begin{aligned} I(R_g, R_b; s_g, s_b) &= H(R_g) - H(R_g | s_g) + H(R_b) - H(R_b | s_b) \\ &= \frac{1}{2}\log(2\pi e \mathrm{Var}(R_g)) - \frac{1}{2}\log(2\pi e \mathrm{Var}(R_g | s_g)) + \\ &+ \frac{1}{2}\log(2\pi e \mathrm{Var}(R_b)) - \frac{1}{2}\log(2\pi e \mathrm{Var}(R_b | s_b)) \\ &= \frac{1}{2}\log\left(\frac{\mathrm{Var}(R_g)}{\mathrm{Var}(R_g | s_g)}\right) + \frac{1}{2}\log\left(\frac{\mathrm{Var}(R_b)}{\mathrm{Var}(R_b | s_b)}\right). \end{aligned} \tag{A.2}$$

Since (R_g, s_g) and (R_b, s_b) are jointly Gaussian, conditional variances satisfy

$$\mathrm{Var}(R_g | s_g) = \mathrm{Var}(R_g)(1 - \rho^2_{R_g, s_g}), \quad \mathrm{Var}(R_b | s_b) = \mathrm{Var}(R_b)(1 - \rho^2_{R_b, s_b}) \tag{A.3}$$

where ρ_{R_g, s_g} and ρ_{R_b, s_b} denote the corresponding correlation coefficient. From the signal structure and the independence assumption, we obtain

$$\rho^2_{R_g, s_g} = \frac{\mathrm{Cov}^2(R_g, s_g)}{\mathrm{Var}(R_g)\mathrm{Var}(s_g)} = \frac{\alpha_g^2 \mathrm{Var}(R_g)}{\mathrm{Var}(s_g)} = \frac{\alpha_g^2 \sigma^2_{R_g}}{\alpha_g^2 \sigma^2_{R_g} + \sigma^2_g}. \tag{A.4}$$

An analogous expression holds for $\rho^2_{R_b,s_b}$. Gathering everything together and plugging it into equation (A.2), we obtain equation (10), as claimed.

□

B First-Order Conditions

Proposition B.1. *Given the optimization problem* (12) *subject to the constraints* (13) *and* (14)*, the first-order conditions (FOC) are given in* (16)*.*

Proof. Given the Lagrangian (15), we need to compute

$$\begin{cases} \frac{\partial \mathscr{L}}{\partial c_1} = 0, \\ \frac{\partial \mathscr{L}}{\partial \alpha_g} = 0, \\ \frac{\partial \mathscr{L}}{\partial \alpha_b} = 0, \\ \frac{\partial \mathscr{L}}{\partial \alpha_f} = 0, \\ \frac{\partial \mathscr{L}}{\partial \lambda_1} = 0, \\ \frac{\partial \mathscr{L}}{\partial \lambda_2} = 0, \end{cases} \tag{B.1}$$

which leads to

$$\begin{cases} \mathbb{E}[\frac{\partial u(c_1,E)}{\partial c_1} + \beta \frac{\partial u(c_2,E)}{\partial c_1}] - \lambda_1 \frac{\partial I(R_g,R_b;s_g,s_b)}{\partial c_1} = 0, \\ \mathbb{E}[\frac{\partial u(c_1,E)}{\partial \alpha_g} + \beta \frac{\partial u(c_2,E)}{\partial \alpha_g}] - \lambda_1 \frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_g} - \lambda_2 = 0, \\ \mathbb{E}[\frac{\partial u(c_1,E)}{\partial \alpha_b} + \beta \frac{\partial u(c_2,E)}{\partial \alpha_b}] - \lambda_1 \frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_b} - \lambda_2 = 0, \\ \mathbb{E}[\frac{\partial u(c_1,E)}{\partial \alpha_f} + \beta \frac{\partial u(c_2,E)}{\partial \alpha_f}] - \lambda_1 \frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_f} - \lambda_2 = 0, \\ I(R_g,R_b;s_g,s_b) - \kappa = 0, \\ \alpha_g + \alpha_b + \alpha_f - 1 = 0. \end{cases} \tag{B.2}$$

Using the chain rule, we can see that

$$\begin{aligned} \frac{\partial u(c_1,E)}{\partial c_1} &= c_1^{-\gamma} \\ \frac{\partial u(c_2,E)}{\partial c_1} &= -c_2^{-\gamma} R_p, \end{aligned}$$

and that

$$\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial c_1} = 0.$$

Let us now compute $\frac{\partial u(c_i,E)}{\partial \alpha_g}$, $\frac{\partial u(c_i,E)}{\partial \alpha_b}$ and $\frac{\partial u(c_i,E)}{\partial \alpha_f}$, for $i = 1,2$ and $\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_g}$, $\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_b}$ and $\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_f}$. We have

$$\frac{\partial u(c_1,E)}{\partial \alpha_g} = \frac{\partial u(c_1,E)}{\partial \alpha_b} = \frac{\partial u(c_1,E)}{\partial \alpha_f} = 0$$

and

$$\frac{\partial u(c_2,E)}{\partial \alpha_g} = c_2^{-\gamma}\frac{\partial c_2}{\partial \alpha_g} = c_2^{-\gamma}(e_1-c_1)\frac{\partial R_p}{\partial \alpha_g} = c_2^{-\gamma}(e_1-c_1)R_g,$$
$$\frac{\partial u(c_2,E)}{\partial \alpha_b} = c_2^{-\gamma}\frac{\partial c_2}{\partial \alpha_b} - \phi E_b = c_2^{-\gamma}(e_1-c_1)\frac{\partial R_p}{\partial \alpha_b} - \phi E_b = c_2^{-\gamma}(e_1-c_1)R_b - \phi E_b,$$
$$\frac{\partial u(c_2,E)}{\partial \alpha_f} = c_2^{-\gamma}\frac{\partial c_2}{\partial \alpha_f} = c_2^{-\gamma}(e_1-c_1)\frac{\partial R_p}{\partial \alpha_f} = c_2^{-\gamma}(e_1-c_1)R_f.$$

We notice that

$$\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_g} = \frac{\partial}{\partial \alpha_g}\left(\frac{1}{2}\log\left(\frac{\mathrm{Var}(R_g)}{\mathrm{Var}(R_g|s_g)}\right)\right) = \frac{\partial}{\partial \alpha_g}\left(\frac{1}{2}\log\left(\frac{\sigma^2_{R_g}}{\frac{\sigma^2_{R_g}}{(1+\alpha_g^2\sigma^2_{R_g}/\sigma^2_g)}}\right)\right) = \tag{B.3}$$
$$= \frac{\partial}{\partial \alpha_g}\left(\frac{1}{2}\log\left(1+\frac{\alpha_g^2\sigma^2_{R_g}}{\sigma^2_g}\right)\right) = \frac{\alpha_g\sigma^2_{R_g}}{\sigma^2_g+\alpha_g^2\sigma^2_{R_g}},$$
$$\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_b} = \frac{\partial}{\partial \alpha_b}\left(\frac{1}{2}\log\left(\frac{\mathrm{Var}(R_b)}{\mathrm{Var}(R_b|s_b)}\right)\right) = \frac{\partial}{\partial \alpha_b}\left(\frac{1}{2}\log\left(\frac{\sigma^2_{R_b}}{\frac{\sigma^2_{R_b}}{(1+\alpha_b^2\sigma^2_{R_b}/\sigma^2_b)}}\right)\right) = \tag{B.4}$$
$$= \frac{\partial}{\partial \alpha_b}\left(\frac{1}{2}\log\left(1+\frac{\alpha_b^2\sigma^2_{R_b}}{\sigma^2_b}\right)\right) = \frac{\alpha_b\sigma^2_{R_b}}{\sigma^2_b+\alpha_b^2\sigma^2_{R_b}},$$
$$\frac{\partial I(R_g,R_b;s_g,s_b)}{\partial \alpha_f} = 0, \tag{B.5}$$

where in (B.3) and (B.4) we used the expressions in (A.3) and in (A.4). Plugging in all together in (B.2) and taking non-stochastic terms outside the expectation yields (16), as claimed. □

C Case where the initial consumption is fixed

In this Appendix we examine a variant of the problem discussed in Section (4), in which the initial consumption c_1 is fixed rather than optimized. Specifically, at time $t=1$ the investor receives the stochastic amount e_1 and consumes a predetermined amount c_1, with $0 \le c_1 \le e_1$, while the remaining wealth $e_1 - c_1$ is allocated across the three assets. Following the literature, a common benchmark is to set c_1 equal to 70% of the initial amount e_1. At time $t=2$, the investor receives the stochastic amount e_2 and consumes everything. The consumption c_2 satisfies the usual constraint (11), while all other assumptions from the main model remain unchanged.

The problem can therefore be written as:

$$\max_{\alpha_g,\alpha_b,\alpha_f} \mathbb{E}[u(c_1,E)+\beta u(c_2,E)] \tag{C.1}$$

subject to constraints (13) and (14) with $\alpha_g \ge 0,\ \alpha_b \ge 0,\ \alpha_f \ge 0$.

The associated Lagrangian reads:

$$\mathscr{L} = \mathbb{E}[u(c_1,E) + \beta u(c_2,E)] - \lambda_1 \left(I(R_g,R_b;s_g,s_b) - \kappa\right) - \lambda_2(\alpha_g + \alpha_b + \alpha_f - 1),$$

where $\mathscr{L} := \mathscr{L}(\alpha_g,\alpha_b,\alpha_f,\lambda_1,\lambda_2)$, and where λ_1,λ_2 are the Lagrange multipliers.

The first order conditions can be written as

$$\begin{cases} \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1)R_g] - \lambda_1 \frac{\alpha_g \sigma^2_{R_g}}{\sigma_g^2 + \alpha_g^2 \sigma^2_{R_g}} - \lambda_2 = 0, \\ \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1)R_b - \beta\phi E_b] - \lambda_1 \frac{\alpha_b \sigma^2_{R_b}}{\sigma_b^2 + \alpha_b^2 \sigma^2_{R_b}} - \lambda_2 = 0, \\ \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1)R_f] - \lambda_2 = 0, \\ I(R_g,R_b;s_g,s_b) - \kappa = 0, \\ \alpha_g + \alpha_b + \alpha_f - 1 = 0, \end{cases} \tag{C.2}$$

where

$$I(R_g,R_b;s_g,s_b) = \frac{1}{2}\log\left(1 + \frac{\alpha_g^2 \sigma^2_{R_g}}{\sigma_g^2}\right) + \frac{1}{2}\log\left(1 + \frac{\alpha_b^2 \sigma^2_{R_b}}{\sigma_b^2}\right),$$

and where

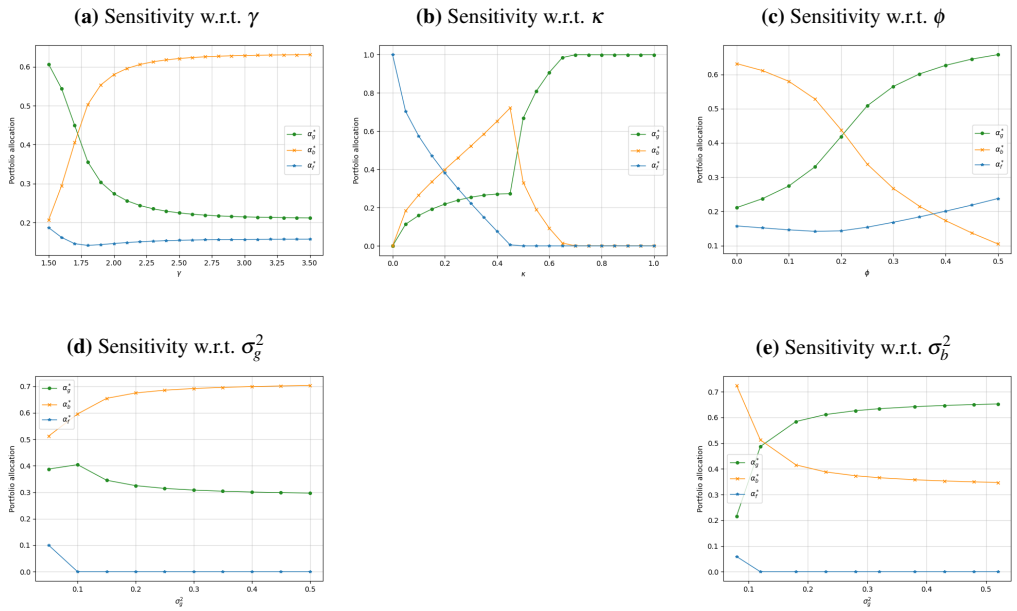
$$c_2 = (e_1 - c_1)(\alpha_g R_g + \alpha_b R_b + \alpha_f R_f) + e_2,$$

with $\alpha_g \geq 0,\ \alpha_b \geq 0,\ \alpha_f \geq 0$.

Assuming the parameters' specification given in Table 1 and fixing $c_1 = 0.7e_1$, we obtain the optimal allocations:

$$\alpha_g^* \approx 0.27, \quad \alpha_b^* \approx 0.58, \quad \alpha_f^* \approx 0.15. \tag{C.3}$$

Figure C.1: Optimal Allocation Responses to Parameter Changes



Notes: This figure show the responses of $\alpha_g^*, \alpha_b^*, \alpha_f^*$ to variation in parameters γ (panel a), κ (panel b), ϕ (panel c), σ_g^2 (panel d), σ_b^2 (panel e) when consumption is fixed.

D Case with non-linear environmental impact

In this Appendix, we analyze a variant of the model presented in Section (4), in which the investor's utility from consumption and environmental impact takes a quadratic form in the environmental component. Specifically, we assume that the iso-elastic (CRRA) utility function is given by

$$u(c,E) = \frac{c^{1-\gamma}}{1-\gamma} - \phi E^2 \tag{D.1}$$

$$= \frac{c^{1-\gamma}}{1-\gamma} - \phi \alpha_b^2 E_b^2. \tag{D.2}$$

All other assumptions of the model remain unchanged.

We therefore study the problem

$$\max_{c_1,\alpha_g,\alpha_b,\alpha_f} \mathbb{E}[u(c_1,E) + \beta u(c_2,E)] \tag{D.3}$$

subject to constraints (13) and (14) with $\alpha_g \geq 0,\ \alpha_b \geq 0,\ \alpha_f \geq 0$, and with $0 \leq c_1 \leq e_1$. In this specification the utility functions become $u(c_1,E) = \frac{c_1^{1-\gamma}}{1-\gamma}$ and $u(c_2,E) = \frac{c_2^{1-\gamma}}{1-\gamma} - \phi \alpha_b^2 E_b^2$.

The associated Lagrangian reads:

$$\mathscr{L} = \mathbb{E}[u(c_1,E) + \beta u(c_2,E)] - \lambda_1 \left(I(R_g,R_b;s_g,s_b) - \kappa\right) - \lambda_2(\alpha_g + \alpha_b + \alpha_f - 1),$$

where $\mathscr{L} := \mathscr{L}(c_1, \alpha_g, \alpha_b, \alpha_f, \lambda_1, \lambda_2)$, and where λ_1, λ_2 are the Lagrange multipliers.

The first order conditions can be written as

$$\begin{cases} \mathbb{E}[c_1^{-\gamma} - \beta c_2^{-\gamma} R_p] = 0 \\ \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1) R_g] - \lambda_1 \frac{\alpha_g \sigma_{R_g}^2}{\sigma_g^2 + \alpha_g^2 \sigma_{R_g}^2} - \lambda_2 = 0, \\ \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1) R_b - 2\beta \alpha_b \phi E_b^2] - \lambda_1 \frac{\alpha_b \sigma_{R_b}^2}{\sigma_b^2 + \alpha_b^2 \sigma_{R_b}^2} - \lambda_2 = 0, \\ \mathbb{E}[\beta c_2^{-\gamma}(e_1 - c_1) R_f] - \lambda_2 = 0, \\ I(R_g, R_b; s_g, s_b) - \kappa = 0, \\ \alpha_g + \alpha_b + \alpha_f - 1 = 0, \end{cases} \tag{D.4}$$

where

$$I(R_g, R_b; s_g, s_b) = \frac{1}{2} \log\left(1 + \frac{\alpha_g^2 \sigma_{R_g}^2}{\sigma_g^2}\right) + \frac{1}{2} \log\left(1 + \frac{\alpha_b^2 \sigma_{R_b}^2}{\sigma_b^2}\right),$$

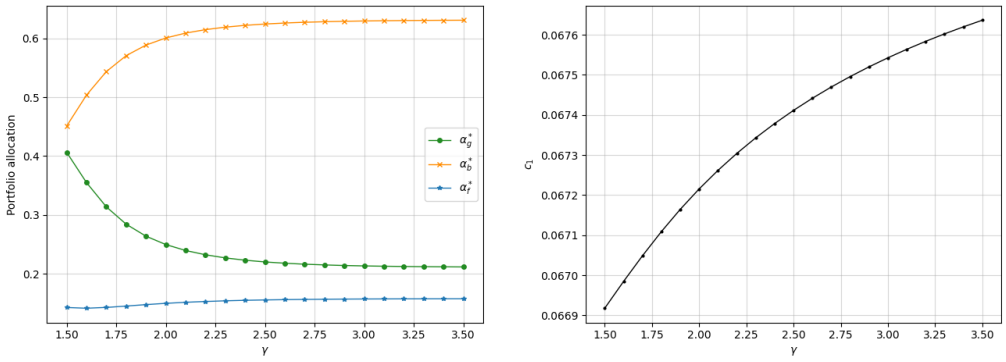
and where

$$c_2 = (e_1 - c_1)(\alpha_g R_g + \alpha_b R_b + \alpha_f R_f) + e_2,$$

with $\alpha_g \geq 0,\ \alpha_b \geq 0,\ \alpha_f \geq 0$, and with $0 \leq c_1 \leq e_1$.
Assuming the parameters' specification given in Table 1, the optimal solution reads:

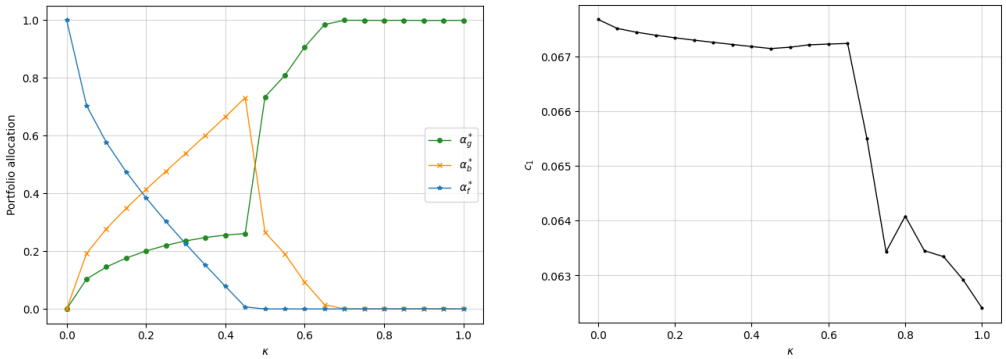
$$c_1^* \approx 0.068, \quad \alpha_g^* \approx 0.25, \quad \alpha_b^* \approx 0.60, \quad \alpha_f^* \approx 0.15. \tag{D.5}$$

Figure D.1: Sensitivity of Portfolio Allocations to Changes in γ



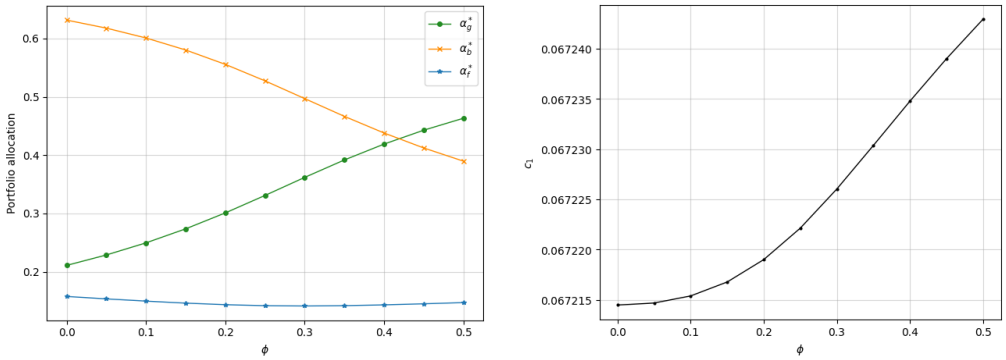
Notes: This figure shows how the three asset-class allocations vary with γ (left panel) and how initial consumption c_1 responds to changes in γ (right panel).

Figure D.2: Sensitivity of Portfolio Allocations to Changes in κ



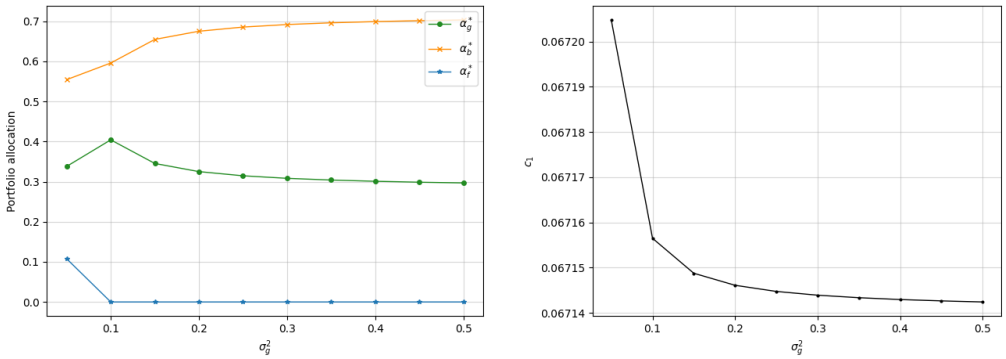
Notes: This figure shows how the three asset-class allocations vary with κ (left panel) and how initial consumption c_1 responds to changes in κ (right panel).

Figure D.3: Sensitivity of Portfolio Allocations to Changes in ϕ



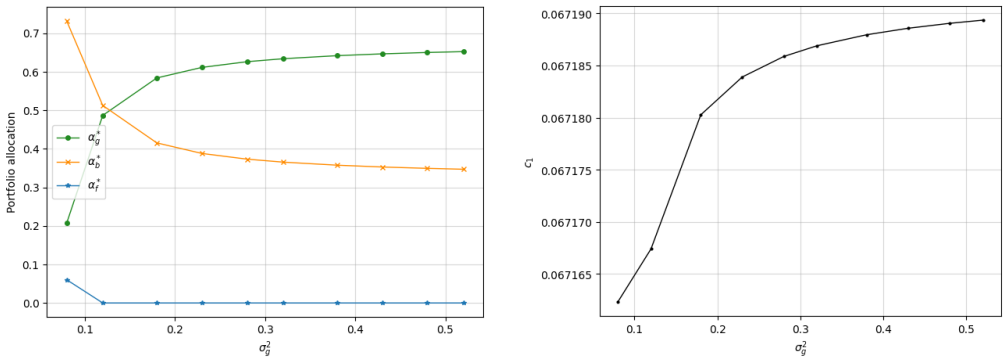
Notes: This figure shows how the three asset-class allocations vary with ϕ (left panel) and how initial consumption c_1 responds to changes in ϕ (right panel).

Figure D.4: Sensitivity of Portfolio Allocations to Changes in σ_g^2



Notes: This figure shows how the three asset-class allocations vary with σ_g^2 (left panel) and how initial consumption c_1 responds to changes in σ_g^2 (right panel).

Figure D.5: Sensitivity of Portfolio Allocations to Changes in σ_b^2



Notes: This figure shows how the three asset-class allocations vary with σ_b^2 (left panel) and how initial consumption c_1 responds to changes in σ_b^2 (right panel).