

Belief Dispersion, Climate Risks, and Returns on Sustainable Investing

Yrjö Koskinen*

Mohammadhossein Lashkaripour[†]

September 20, 2026

*Koskinen (yrjo.koskinen@ucalgary.ca) is the BMO Professor of Sustainable and Transition Finance at the Haskayne School of Business, University of Calgary and a research member with ECGI.

[†]Lashkaripour (mlashkaripour@uvic.ca) is a PostDoc at the University of Victoria.

We thank Jan Starmans, Hany Fahmy, Valeri Sokolovski, Aymeric Bellon, and Michael King, as well as conference and seminar participants at the University of Ottawa, Queen's University, the University of Victoria, Aalto University, the 2025 Canadian Sustainable Finance Network (CSFN) Conference, the 2nd Exeter Sustainable Finance Conference, and the 2026 IFABS Conference for their comments and feedback.

Belief Dispersion, Climate Risks, and Returns on Sustainable Investing

Abstract

Recent climate asset pricing models imply a positive brown-minus-green (BMG) expected return premium as compensation for physical climate and transition risks, yet empirical estimates of that premium yield mixed results. We argue that disagreement about climate damages is a missing state variable that helps explain the disconnect between theory and empirical evidence. We first present suggestive evidence by measuring climate belief dispersion from variation in carbon emissions tilt in institutional asset managers' portfolios. We show that greater dispersion predicts lower next-quarter BMG returns, even after controlling for conventional risk factors and macro-financial conditions. Motivated by this evidence, we develop an overlapping-generations equilibrium model with purely pecuniary investors who take into account climate damages in their portfolio decisions but disagree about their magnitude. Disagreement generates speculative trading, distorts risk sharing, and affects transition risk through belief-driven political pressure over an emissions cap. The model shows that the expected BMG return can be positive, zero, or negative depending on climate belief dispersion and policy pressure; sufficiently convex climate damages can restore a uniformly positive BMG premium.

Keywords: Heterogeneous beliefs; Climate change; Climate risks; Transition risks; Speculation; Sustainable investing

1 Introduction

Climate risk is increasingly viewed as a first-order determinant of long-run cash flows, discount rates, and transition policy. Yet its asset pricing implications remain unsettled. Many climate asset pricing models predict that brown assets should offer higher expected returns than green assets to compensate for their greater exposure to physical climate risk, transition risk, or investor exclusion. Under this view, a brown-minus-green (BMG) portfolio should earn a positive expected return. The empirical evidence is more mixed: some studies estimate that the expected BMG return is positive (Avramov, Lioui, Liu, and Tarelli, 2025), while others suggest that it varies substantially over time and sometimes changes sign (Sautner, Van Lent, Vilkov, and Zhang, 2023). In some settings, BMG returns appear more consistent with mispricing than with a stable risk premium (Atilgan, Demirtas, Edmans, and Gunaydin, 2024). This gap between sharp theoretical predictions and mixed empirical findings motivates our analysis.

We propose that an important missing state variable from current models is dispersion in beliefs about the economic damages of climate change and the policy response to those damages. Investors may agree about firms' current financial fundamentals but disagree about the magnitude of future climate damages, the likelihood of stringent climate policy, and the implications of emissions exposure for future cash flows. Recent studies document noticeable disagreement along several of these dimensions (Ceccarelli and Ramelli, 2024; Giglio, Maggiori, Stroebel, Tan, Utkus, and Xu, 2025). More importantly, unlike disagreement about many financial fundamentals, which can be resolved relatively quickly as observable outcomes are realized, disagreement about climate change may persist for long periods and can even widen over time. In fact, climate damages unfold over long horizons, are realized with considerable noise, and are often interpreted through broader political and economic views. As a result, differences in climate beliefs may take a long time to resolve and can even widen over time. Public opinion data are consistent with this idea. For example, Figure 1 illustrates rising cross-state dispersion in climate beliefs in the United States using the Yale Program on Climate Change Communication's Climate Opinion Maps (YCOM) over 2010–2024.




Figure 1. Belief dispersion across U.S. states from the Yale Climate Opinion Maps (YCOM) estimates. For each year, dispersion is computed as the unweighted standard deviation across states for each belief item (e.g., ‘global warming is happening,’ ‘caused mostly by human activities,’ and attribution/experience items), then averaged across those belief items to form the annual dispersion series.

Our central premise is that disagreement about expected climate damages is priced as it induces speculative trading and alters equilibrium risk sharing. When investors disagree about the severity of climate damages, they take opposing positions on climate-related risks, generating additional components of equilibrium expected returns. These components depend on both (i) the average level of perceived climate damages and (ii) the dispersion of those perceptions across investors. In such an environment, the BMG return premium need not be stable: it can be positive, close to zero, or negative depending on the extent of disagreement and on how climate policy responds to the distribution of beliefs. We examine this mechanism using a tractable equilibrium model.

Before developing the model, we present suggestive empirical evidence supporting our premise that disagreement about climate damages is priced in financial markets. Because investors’ climate beliefs are not directly observable, we infer disagreement from institutional portfolio holdings. For each asset manager, we measure an emissions tilt relative to an industry benchmark. Specifically, we measure firms’ carbon emissions intensity using Scope 1 and Scope 2 emissions scaled by sales, compare each firm’s intensity with that of other firms in the same industry and quarter, and then calculate the manager’s portfolio-weighted exposure to

these industry-adjusted intensities. We then define belief dispersion as the cross-sectional standard deviation of the component of that exposure left unexplained by standard portfolio characteristics, manager type, and investment strategy. The resulting series trends upward over the sample, rising modestly after the December 2015 Paris Agreement and more sharply after 2020. In our regression analysis, we show that a one-standard-deviation increase in climate belief dispersion is followed by roughly a one-percentage-point decline in the next quarter's BMG return, and this negative relation persists after controlling for the Fama–French five factors and macro-financial conditions.

Motivated by these patterns, we develop an overlapping-generations equilibrium model with two types of investors with purely pecuniary motives who disagree about the expected economic damages caused by climate change. Dividends are produced from energy and capital, and carbon emissions arise from energy use. Emissions accumulate over time and raise temperature, which destroys wealth through a climate-damage function. Investors share common beliefs about conventional financial fundamentals but disagree about the expected economic damages caused by climate change. Pessimists overweight high-temperature states where damages are larger; optimists overweight low-temperature states where damages are smaller. Because asset prices are forward-looking, if investors anticipate climate damages, stock prices adjust to reflect this expectation. As a result, price-dividend ratios vary with cumulative emissions, the distribution of beliefs, and the expected climate policy, which is represented through a carbon emissions cap.

Because both investor types take into account climate damages but evaluate them under different subjective probability measures, disagreement generates three effects. First, the two investor types assign different expected damages to the same emissions exposure, so they willingly take opposite sides of climate-related risks and trade speculatively. Second, each type's portfolio is no longer a pure risk–return trade-off but a best response to the other type's holdings. In fact, each type views the opposing position as suboptimal under its own beliefs and therefore adjusts its portfolio accordingly, which distorts risk sharing and generates equilibrium hedging demand. Third, because the state of climate change affects both dividends and valuation ratios, differences in climate beliefs also translate into differences in perceived asset returns. Each of these effects is associated with a distinct pricing component. Equilibrium expected returns can therefore be represented by a standard two-factor climate asset pricing model with financial and climate risk factors, augmented by three climate belief-driven components: speculation, hedging, and a belief-driven valuation wedge.

The interaction between these two risk factors and the three climate belief-driven pricing components de-

livers two main results: First, belief dispersion can make the BMG expected return positive, zero, or negative, depending on how far each type's belief deviates from the objective belief and on each type's population share. This reconciles the standard prediction of a positive brown premium with the mixed empirical evidence on realized brown-minus-green returns. Second, sufficiently convex climate damages can restore a positive BMG premium because, under such conditions, even optimistic investors find marginal warming costly enough to discount brown assets more heavily.

We then extend the model to examine how the political mechanism shapes the dynamics of the BMG premium. We first consider a probabilistic-voting model under which climate policy responds smoothly to belief-driven political pressure: as the share of pessimistic investors rises, pressure for a tighter emissions cap increases. When belief shares endogenously evolve with past realized returns, this creates a belief-policy feedback channel. Strong realized BMG performance improves the relative performance of climate optimists, increases their share in the next entering cohort of investors, weakens political pressure for stringent climate policy, and ultimately relaxes the emissions cap. Two opposing forces follow. On the one hand, greater optimism increases investors' willingness to hold brown assets, reducing the compensation they require (*demand effect*). On the other hand, a looser cap increases the brown market weight investors must collectively hold, raising the compensation needed to absorb that exposure (*cap effect*). The resulting BMG dynamics therefore depend on the relative strength of these two forces: the model generates mean reversion when the demand effect dominates the cap effect, and momentum when the cap effect is stronger.

Second, we study an extension of the political process with median voting. This extension shows that the policy-feedback mechanism can be discontinuous rather than smooth. Under median voting, the emissions cap is determined by the preferred policy of the median investor: if pessimists form the majority, the economy adopts a stringent emissions cap; if optimists form the majority, it adopts a more lenient cap. Holding each type's climate expectations fixed, changes in belief shares affect the cap only when they change the majority type, although investor demand can adjust within each regime. An increase in optimism can therefore move the BMG premium before it moves policy. Once optimists gain a majority, the cap relaxes and the brown market weight rises, but it also changes the exposure of asset returns to climate shocks, since a more lenient cap translates into larger expected climate damages. As a result, under median voting, changes in beliefs—and hence in political pressure—initially have only a gradual effect on the BMG premium, which then shifts sharply once the median type changes. This implies that a median-voting economy is associated with higher

climate-related uncertainty than a probabilistic-voting economy.

Related Literature. This paper contributes to two strands of literature. First, it contributes to climate asset pricing. One branch of this literature derives a positive BMG premium from non-pecuniary preferences: investors who value greenness bid up green assets and require compensation for holding brown ones (Pástor, Stambaugh, and Taylor, 2021; Pedersen, Fitzgibbons, and Pomorski, 2021; Zerbib, 2022). Other studies also imply a positive BMG when green investors exclude brown firms from their opportunity set (Heinkel, Kraus, and Zechner, 2001) or when preferences for green consumption lower green assets’ systematic risk (Albuquerque, Koskinen, and Zhang, 2019). A second branch reaches the same prediction from fundamentals rather than preferences: because warming destroys output and welfare (Weitzman, 2012; Hong, Wang, and Yang, 2023; Barnett, Brock, and Hansen, 2020; Hambel, Kraft, and Schwartz, 2021), investors discount (or demand compensation for holding) claims on emission-intensive production (Ilhan, Sautner, and Vilkov, 2021; Bolton and Kacperczyk, 2021). Anticipated policy responses—including emissions caps, carbon taxes, and other forms of regulation—introduce an additional transition-risk channel that further raises the required return on brown assets (Barnett, 2024; Hsu, Li, and Tsou, 2023).

The empirical evidence, however, does not point to a stable BMG premium. Pástor, Stambaugh, and Taylor (2022) document higher realized returns for green assets than the model in Pástor et al. (2021) predicts; Sautner et al. (2023) find that the price of climate exposure varies over time and across risk types; Bolton and Kacperczyk (2023) find transition risk priced while Zhang (2025) reaches the opposite conclusion on the same data; and Jagannathan, Meier, and Sokolovski (2025) show that the apparent premium largely reflects the concentration of heavy emitters in value portfolios. We show that unstable return spread between brown and green assets need not imply that climate risk is unpriced: time variation and sign reversals arise endogenously when investors disagree about damages, when transition policy responds to the distribution of beliefs, and when valuations capitalize both.

Four studies in this strand are closely related to ours and warrant further discussion. First, De Angelis, Tankov, and Zerbib (2023) analyze an economy with two types of investors—green and traditional—where green investors deviate from the true probability measure by associating dividend outcomes with emissions profiles, while traditional investors focus solely on monetary payoffs. Although their framework incorporates belief heterogeneity, our approach differs in three important respects. First, we do not assume ex ante differences in preferences or beliefs about emissions profiles; instead, beliefs evolve endogenously from observing

the previous welfare damages of climate change. Second, all investors in our model are purely pecuniary. Third, investors share common beliefs about financial performance but disagree only about climate-related damages. In addition, our model endogenizes transition risk.

Second, Baldauf, Garlappi, and Yannelis (2020) examine how climate beliefs affect real estate prices and show that climate risk is capitalized into housing values only in regions where local populations believe in climate change. Our study differs in two key ways. First, we focus on equity markets, where production decisions generate emissions that influence future output. Second, we demonstrate that not only the level of climate beliefs but also their dispersion plays a central role in shaping asset prices and expected return differentials. Relatedly, Ardia, Bluteau, Boudt, and Inghelbrecht (2023) study how the level of climate concerns—proxied by media attention—affects asset prices and the return spread between brown and green assets. We complement this perspective by examining the pricing implications of disagreement among investors alongside their average climate beliefs.

Finally, Garleanu and Pedersen (2025) develop an environmental macro-finance model in which the required return on brown assets relative to green assets depends on the interaction between climate and policy shocks. Similar to our model, their framework implies that the BMG premium can vary over time and change sign. Our mechanism for generating this time variation, however, is different. Besides a different mechanism, a key contribution of our model relative to theirs is to endogenize the link between physical and transition risks, which is central to their analysis but is taken as exogenous. In our framework, climate-belief dispersion affects political pressure over climate policy, so beliefs about future climate damages endogenously shape transition risk.

The second strand to which our paper contributes is the literature on the asset pricing implications of heterogeneous beliefs. Basak (2005), David (2008), and Dumas, Kurshev, and Uppal (2009) establish that investors who interpret common information differently trade speculatively, and that disagreement is priced and generates excess volatility; discrete-time treatments closer to ours include Jouini and Napp (2006) and Hong and Sraer (2016). We apply belief dispersion to climate risk and, departing from this literature, link disagreement to endogenous policy formation. The median-voting extension makes this mechanism more pronounced by introducing a threshold in the policy response. As beliefs shift, policy remains unchanged until the identity of the median voter changes. Once that happens, the emissions cap switches between a stringent and a lenient regime, producing a discrete change in the BMG premium.

The remainder of the paper is organized as follows. Section 2 presents suggestive evidence on belief dispersion and BMG returns. Section 3 develops the theoretical model. Section 4 characterizes equilibrium asset prices and the BMG premium with and without endogenous climate policy. Section 5 concludes.

2 Suggestive Evidence

Our core argument motivating the model is that dispersion in investors' beliefs about climate damages is an important state variable for the brown–green return spread. To provide empirical support for this argument, we need a measure of climate disagreement among investors. Because climate beliefs are not directly observable, we construct a proxy from institutional portfolio holdings and examine whether it predicts subsequent brown-minus-green (BMG) returns.

Our use of institutional holdings is based on revealed portfolio choices. If a manager expects climate damages or the associated policy response to materially reduce a firm's future cash flows or valuation, that expectation should, all else equal, reduce the manager's exposure to emission-intensive firms. We therefore first measure each manager's emissions tilt. We then remove the components of that tilt attributable to observable portfolio characteristics, manager type, and investment strategy, and use the cross-sectional dispersion of the residual component as a proxy for climate-belief dispersion.

This exercise does not require all residual variation in emissions tilts to arise exclusively from climate beliefs. Instead, it requires that residual dispersion co-move with the underlying dispersion in climate beliefs and that it not predict future BMG returns through channels unrelated to the mechanism developed in the model. The latter condition cannot be verified directly. We therefore interpret the empirical evidence as suggestive and use it to motivate the theoretical analysis rather than as a definitive test of the model.

2.1 Data

Quarterly institutional holdings are obtained from Thomson Reuters 13F filings and merged with CRSP security returns and market capitalization, ISS Scope 1 and Scope 2 carbon emissions, Compustat sales, and institutional-manager classifications. The final sample covers from 2014Q1 through 2023Q3, yielding 39 quarterly observations.

The sample period is determined by the availability of ISS emissions data, which spans the years 2012–2023. Emissions are generally disclosed with a substantial delay: annual emissions for a given fiscal year typically become publicly available only during the following calendar year. To avoid assigning managers information that was unavailable when their portfolios were formed, we match each quarter’s holdings to emissions from the most recent fiscal year that would have been disclosed at that time. Operationally, we use emissions from the preceding fiscal year and refer to this matching convention as the one-year emissions lag.

The lagged construction further restricts the beginning of the sample. Holdings in 2013 must be matched to fiscal-year 2012 emissions, but the 2012 ISS file contains only 30 successfully linked firms. After imposing the portfolio-coverage and manager-information requirements described below, the resulting 2013 dispersion measure would be based on only one or two managers per quarter and would therefore be uninformative. From 2014 onward, the lagged-emissions universe contains more than 2,200 firms and supports approximately 2,500 to 4,000 institutional managers per quarter.

2.2 Climate Beliefs

For firm i , we measure emissions intensity as Scope 1 plus Scope 2 emissions divided by trailing four-quarter sales, using the one-year emissions lag described above. Scaling emissions by sales helps account for firm size, because emission levels rise mechanically with firm size, so a level-based tilt would load on a size factor. Sales are trailing for the same information-availability reason that motivates the emissions lag. We exclude Scope 3, whose coverage is sparse over our sample and which double counts emissions along the supply chain.

We transform intensity by $\log(1 + x)$, winsorize it within Fama–French 12 industry–quarter cells, and subtract the industry–quarter mean. The purpose of this log transformation is to reduce the influence of the highly skewed emissions distribution, which would otherwise allow a small number of utilities and materials firms to dominate portfolio-level measures. More importantly, we adjust emissions intensity within industries. The reason is that much of the variation in emissions intensity comes from differences across industries, meaning that an unadjusted measure would largely capture sector allocation rather than managers’ views about climate exposure. By comparing firms within the same industry and quarter, we focus instead on the manager’s choice among firms facing broadly similar production technologies. This also makes our measure consistent with the construction of the BMG portfolio, which likewise sorts firms within industries.

A manager’s climate tilt is measured by the value-weighted average of industry-adjusted exposure across the firms in its portfolio with valid emissions data. So, a tilt of zero corresponds to industry-average carbon exposure. Because ISS does not cover the full 13F universe, this average is taken over part of the portfolio that has available emissions data. Specifically, we define a manager’s *emissions coverage* as the share of reported equity value held in firms with matched lagged emissions, and the preferred sample retains manager-quarters with coverage of at least 70%. In calculating tilt, we renormalize the portfolio weights of covered firms to sum to one; thus, the construction implicitly assumes that the covered holdings are informative about the manager’s overall climate positioning. This assumption becomes less credible as coverage declines. In particular, a tilt calculated from a small and potentially unrepresentative portion of the portfolio is measured with greater error. Because D_t is the cross-sectional standard deviation of residual manager tilts, such manager-specific measurement error may inflate measured dispersion rather than average out. Moreover, because emissions coverage improves over the sample, including managers with low coverage could generate a spurious downward trend in dispersion as measurement quality improves. With all that said, we assess the sensitivity of the results to a less restrictive 50% coverage threshold in the Appendix.

We residualize manager-level tilt on portfolio size, turnover, concentration, permanent-QIX passive status, manager type, and quarter fixed effects. Size and concentration remove a mechanical effect: diversified portfolios are pulled toward the market tilt of zero, while concentrated portfolios show extreme tilts for reasons unrelated to climate. Turnover proxies for horizon. Passive status and manager type remove exposure that is inherited rather than chosen, since a quasi-indexer’s carbon tilt largely reflects its benchmark and manager types differ in regulatory and fiduciary constraints even when beliefs coincide.

Quarter fixed effects deserve emphasis, because they are what makes this a measure of disagreement rather than of the average level of climate concern. They absorb all common time variation in the average tilt, including movements in institutions’ average climate concern. What remains is purely cross-sectional, and our proxy

$$D_t \equiv \text{sd}_m(\hat{\epsilon}_{m,t}) \tag{1}$$

is the cross-sectional standard deviation of residual tilt $\hat{\epsilon}_{m,t}$.

Figure 2 plots D_t . Dispersion is relatively low around the December 2015 Paris Agreement and subsequently trends upward: its average is 0.420 through 2015Q4, 0.455 thereafter, and 0.503 over the final eight

quarters. The increase is not monotonic, and the Agreement is a reference point rather than an identifying event. The pattern is consistent with climate exposure becoming a more contested dimension of portfolio choice as transition policy becomes more salient.




Figure 2. Residual institutional climate-belief dispersion. The series is the quarterly cross-sectional standard deviation of residual continuous emissions tilt among managers with at least 70% emissions coverage. The dashed vertical line marks the adoption of the Paris Agreement in December 2015. The sample begins in 2014Q1; 2013 is excluded because lagged emissions coverage leaves only one or two usable managers per quarter.

2.2.1 External Validity

A proxy of this kind should meet two requirements: it should move with independent measures of the same phenomenon, and it should be distinguishable from related constructs with which it might be confused. We report both tests in the Appendix and summarize them here.

The first regresses the news-based climate-policy uncertainty indexes of Gavriilidis, Känzig, Raghavan, and Stock (2026), averaged within quarter, on our dispersion measure, with both series standardized so that the slope reads as a standard-deviation response (equation (A.1)). The two measures share no input data and no methodology—one is built from newspaper text, the other from portfolio holdings—so common measurement error is implausible, and a positive slope indicates that institutional carbon positions are most dispersed when climate policy is most contested in public debate. Table 4 reports slopes of 0.57–0.65 across the narrow, broad-dictionary, and legacy indexes, significant in all three.

The second test asks whether D_t merely proxies for generalized uncertainty, in which case it would amount to a repackaged measure of macroeconomic uncertainty. Equation (A.2) replaces the dependent variable with the broad U.S. Economic Policy Uncertainty Index of Baker, Bloom, and Davis (2016) and adds climate-policy uncertainty as a control. If dispersed emissions tilt reflects economy-wide uncertainty, dispersion should retain explanatory power for this placebo outcome even after conditioning on the climate component. If it is specific to climate-related uncertainty, it should not. Table 5 supports the latter interpretation: the unconditional relationship is weak, and conditioning on either climate index renders the dispersion coefficient insignificant, while climate-policy uncertainty remains strongly related to the placebo outcome. Thus, to the extent that D_t captures a component of aggregate uncertainty, it appears to capture the climate-policy component. Because the short sample limits statistical power, we view this result as a weak rather than decisive rejection.

2.3 BMG Return and the Dispersion in Climate Beliefs

The BMG portfolio is constructed from the broader CRSP–ISS common-stock universe rather than being restricted to stocks held by the institutional managers in our sample. In each industry and quarter, we rank firms by lagged emissions intensity and classify those in the top quartile as brown and those in the bottom quartile as green. We value weight both portfolios using lagged market capitalization and compound monthly returns to obtain quarterly BMG returns. We then estimate

$$R_{t+1}^{BMG} = \alpha + \beta D_t + \Gamma' F_{t+1} + \Delta' Z_t + \varepsilon_{t+1}, \quad (2)$$

where F_{t+1} contains the Fama–French five factors and Z_t contains log VIX and log Partisan Conflict Index. To make magnitudes comparable, we standardize disagreement measures to zero mean and unit variance within the estimation sample. Coefficients therefore measure the percentage-point change in next-quarter BMG returns associated with a one-standard-deviation change in the focal variable. Newey–West inference uses four quarterly lags, and the tables report t -statistics in parentheses.

Table 1 presents the main result. Residual dispersion enters with a negative coefficient in all three specifications. Without controls, a one-standard-deviation increase in dispersion is associated with a 0.86 percentage-point lower BMG return in the following quarter. The estimate is 0.76 percentage points after adding the Fama–French five factors and 1.15 percentage points in the fully controlled specification that also includes VIX

and the partisan conflict index. The first two estimates are significant at the 10% level, while the fully controlled estimate is significant at the 1% level. Given that the sample contains only 39 quarters, we also assess the result using circular shifts of the dispersion series relative to future BMG returns. The corresponding p -values range from 0.077 to 0.128. Overall, the negative relationship is fairly consistent across specifications, and the estimated magnitudes remain economically meaningful.

Table 1. Residual institutional climate-belief dispersion and future BMG returns

	(1)	(2)	(3)
Residual dispersion	-0.859*	-0.759*	-1.151***
	(-1.85)	(-1.89)	(-3.05)
MKT–RF		-0.054	-0.058
		(-1.05)	(-1.30)
SMB		-0.137	-0.269***
		(-1.60)	(-2.67)
HML		0.207**	0.277***
		(2.52)	(2.81)
RMW		-0.257	-0.322**
		(-1.36)	(-2.02)
CMA		0.164	0.095
		(0.86)	(0.48)
log(VIX)			4.981**
			(2.13)
log(PCI)			2.375
			(1.34)
Constant	-1.411***	-0.913**	-26.87*
	(-2.94)	(-2.06)	(-1.88)
Observations	39	39	39
R^2	0.062	0.442	0.519

Notes: The dependent variable is the signed BMG return in quarter $t + 1$. Disagreement measures are standardized to zero mean and unit standard deviation within the coverage sample. Newey–West t -statistics with four quarterly lags are in parentheses. FF5 controls are measured in $t + 1$; log VIX and log PCI are measured in t . All columns use the same sample.

***, **, and * denote significance at the 1%, 5%, and 10% levels.

Overall, the evidence indicates that periods of greater institutional disagreement about emissions exposure are followed by lower BMG returns. This negative predictive relationship is consistent with the model’s benchmark mean-reversion mechanism, which we develop in the following sections.

3 Model

This section outlines the model structure used to study how heterogeneous beliefs about climate damages and endogenous transition policy shape equilibrium asset prices and the expected return spread between brown and green assets. We proceed as follows. First, we specify the economic environment and asset structure. Second, we model climate change as the accumulation of emissions generated in the dividend-production process and introduce a damage function. Third, we solve the investor’s portfolio problem. Fourth, we characterize belief dispersion regarding climate damages and explain how these beliefs spill over into beliefs about returns. Fifth, we endogenize the share of investors holding optimistic and pessimistic climate beliefs. Sixth, we solve for the equilibrium in the absence of climate policy. Seventh, we introduce climate policy in the form of an emissions cap that responds to belief-driven political pressure, with this pressure determined by the endogenous distribution of beliefs. Finally, we solve for equilibrium expected returns under climate policy.

Notation. Throughout the paper, unless otherwise stated, lowercase letters denote the natural logarithms of their uppercase counterparts whenever both are defined. Boldface letters denote vectors.

3.1 The Economy

Time is discrete and infinite, indexed by $t = 0, 1, 2, \dots$. The economy features overlapping generations (OLG) of investors, with each generation living for two periods. The OLG structure serves two main purposes. First, it introduces a dynamic framework without adding unnecessary technical complexity. Second and more importantly, it allows for persistent belief dispersion over time. In the absence of overlapping generations, investors would learn from others’ behavior over time, leading beliefs to converge in the long run. By contrast, the OLG structure sustains persistent disagreement.¹

¹ Alternatively, we could assume dogmatic beliefs in a multiperiod economy. We prefer the OLG framework, however, because we later endogenize the shares of climate optimists and pessimists based on their relative past performance, and we view this approach as more realistic in the context of climate-change beliefs. Some multiperiod and continuous-time models with dogmatic beliefs allow the distribution of belief types to evolve endogenously over time, as in Atmaz and Basak (2018). However, such models do not explicitly

We define generation t as a cohort of investors born at time t (“young”) and exiting the market at time $t + 1$ (“old”). Each generation contains two types of purely pecuniary investors who disagree about the economic damages caused by climate change. To isolate this source of heterogeneity, we assume that both types share common beliefs about conventional financial fundamentals, such as productivity and capital shocks, but evaluate climate-related outcomes under different subjective probability measures. Pessimists place greater probability weight on high-temperature states, whereas optimists place greater weight on low-temperature states. We index optimists and pessimists by $i = O$ and $i = P$, respectively, and π_t^i denote the population share of type i in generation t . There is no information asymmetry: both types observe the same information and each other’s actions but adhere to their own climate beliefs.

In generation t , an investor of type i is endowed with wealth W_t^i . When young, the investor allocates this wealth across the available assets and consumes terminal wealth W_{t+1}^i when old. For tractability, we assume that each generation consists of one representative investor of each type.

There are three assets available for investment: a risk-free asset in zero net supply, with gross return R_f , and two risky assets, green and brown, denoted by $j = G$ and $j = B$, respectively. Both risky assets generate dividends through a Cobb-Douglas technology using energy and a composite capital input (interpretable as human capital, intangible or tangible capital):

$$Div_{t+1}^j = A \left(K_{t+1}^j \right)^\alpha \left(E_{t+1}^j \right)^{1-\alpha} \quad (3)$$

where $0 < \alpha < 1$, Div_t^j denotes the dividend of asset j , A is the total factor productivity, E_{t+1}^j denotes the energy input, and K_{t+1}^j denotes the composite capital input. Shocks to E_{t+1}^j and K_{t+1}^j are assumed to be orthogonal. Energy used for production generates emissions as follows:

$$Emission_{t+1}^j = \psi^j E_{t+1}^j \quad (4)$$

where ψ^j denotes the emission intensity of asset j . We define an asset as brown if it has a higher emission

capture the underlying mechanism or the role of generational differences in climate beliefs, despite substantial evidence that these differences exist. For example, the Pew Research Center reports that, in the United States, younger evangelicals are far more concerned about climate change than their elders. See:

<https://www.pewresearch.org/short-reads/2022/12/07/younger-evangelicals-in-the-u-s-are-more-concerned->

intensity than a green asset:

$$\psi^B > \psi^G \quad (5)$$

The prices of risky assets are denoted by P_t^j , and their gross return is therefore given by $R_{t+1}^j = \frac{P_{t+1}^j + Div_{t+1}^j}{P_t^j}$.

Remark 1: *Because investors do not consume when young and instead maximize the utility of terminal wealth, there is no consumption–investment trade-off within a given generation. At first glance, this may suggest that each generation behaves identically—simply choosing the portfolio that maximizes terminal utility—making the overlapping-generations structure appear superficial. However, since state variables evolve over time, investors in different generations face different opportunity sets and therefore make different optimal choices. In this sense, the OLG framework remains essential. This modeling approach, namely overlapping generations of investors who care only about terminal wealth but operate in a time-varying environment, has been used in other studies such as Acharya and Pedersen (2005).*

3.2 Climate Change

Aggregate emissions from green and brown production accumulate in the atmosphere and drive climate change.

Let X_t^j denote the weight of asset j in the market portfolio in generation t . Thus, aggregate emissions are:

$$Emission_{t+1}^{agg} = X_t^G Emission_{t+1}^G + X_t^B Emission_{t+1}^B \quad (6)$$

Climate science suggests that there exists an approximately linear relationship between cumulative emissions and the average global temperature (Matthews, Gillett, Stott, and Zickfeld, 2009). Accordingly, we model climate change—defined as the change in the Earth’s average surface temperature—as

$$\Delta T_t \approx v(Emission_{t+1}^{agg} + Emission_t^{cum}) \quad (7)$$

where v is known as Carbon-Climate Response (CCR), $Emission_{t+1}^{agg}$ is defined above, and $Emission_t^{cum}$ denotes the cumulative emissions inherited from the past:

$$Emission_t^{cum} = \sum_{s=1}^t Emission_s^{agg} \quad (8)$$

Climate change destroys wealth and reduces terminal output, which investors consume before exiting the market. We capture this effect through a damage function $D(\Delta T_t)$, which can be interpreted as a discount factor on output (Barnett, 2024). Our framework accommodates a wide range of specifications for $D(\Delta T_t)$, as detailed in the Online Appendix. In the main analysis, we adopt the standard polynomial specification used in Dynamic Integrated Climate–Economy (DICE) models:

$$D(\Delta T_t) = \frac{1}{1 + \delta \Delta T_t^\eta} \quad (9)$$

where δ is the damage intensity parameter and η determines the damage convexity with higher η implies more destructive damages.

Investors differ in their expectations about climate damages $D(\Delta T_t)$. Because ΔT_t is the sole input to the damage function, differences in beliefs about damages arise from heterogeneous expectations about temperature changes. Formally, belief dispersion can be expressed as $E_t^i[\Delta T_t] \neq E_t^{i-}[\Delta T_t]$, where $E_t^i[\cdot]$ denotes the expectation under type i 's belief and $E_t^{i-}[\cdot]$ denotes the expectation under the other type's belief. We later introduce additional structure on belief heterogeneity to study how the extent of disagreement affects the return spread between brown and green assets.

We also later characterize climate policy in Section 4.2 as a cap on aggregate emissions (Cap_t) that is endogenously determined through belief-driven political pressure.

Remark 2: *Although we use a DICE-style damage function, our climate change model is highly stylized compared to DICE and other Integrated Assessment Models (IAMs). For example, we do not consider a multilayer temperature system or carbon diffusion between layers. The main goal of this stylized modelling is to summarize climate change into a single representative variable that allows us to study which asset pricing effects are due to climate risk and which are due to belief dispersion about climate damages. In a multidimensional setting, this task would be difficult, as many interactions would show up in the model, preventing us from deriving an interpretable closed-form solution.*

3.3 Investors Optimization

In each generation both investor types maximize expected utility of terminal wealth, accounting for climate damages:

$$\underset{X_t^{i,G}, X_t^{i,B}}{Max} E_t^i [U(D(\Delta T_t)W_{t+1}^i)] \quad (10)$$

subject to a budget constraint, where $E_t^i[\cdot]$ denotes expectation under type i 's belief.

Valuation ratios and returns

We established earlier the link between the dividend and emissions. Since these emissions can cause damages in the future, the economy accounts for them, and thus prices, in addition to dividends, should respond to emissions. Put differently, prices in this economy are forward-looking: what a claim is worth today depends on expected future dividends and on the compensation investors require for holding it, where both are associated with emissions. The natural device for carrying this dependence is the price–dividend ratio, which links prices to fundamentals, varies with the state of the economy, and is the margin through which equilibrium expected returns adjust. Below, we define valuation ratios, express returns in terms of them, and derive the conditional return moments that the portfolio problem takes as inputs.

We can rewrite the gross return and log return on asset j as follows:

$$R_{t+1}^j = \frac{P_{t+1}^j + Div_{t+1}^j}{P_t^j} = \frac{1 + PD_{t+1}^j}{PD_t^j} \frac{Div_{t+1}^j}{Div_t^j}, \quad r_{t+1}^j = \Delta d_{t+1}^j + \log \left(1 + e^{pd_{t+1}^j} \right) - pd_t^j \quad (11)$$

where PD_t^j and pd_t^j denote the price–dividend and log price–dividend ratio, respectively, and $\Delta d_{t+1}^j = \log Div_{t+1}^j - \log Div_t^j$ denotes log dividend growth. Following Campbell and Shiller (1988), we log-linearize (11) around the unconditional mean $\overline{pd}^j$:

$$r_{t+1}^j \approx k_j + \rho_j pd_{t+1}^j - pd_t^j + \Delta d_{t+1}^j, \quad \rho_j = \frac{e^{\overline{pd}^j}}{1 + e^{\overline{pd}^j}} \in (0, 1) \quad (12)$$

with $k_j = \log(1 + e^{\overline{pd}^j}) - \rho_j \overline{pd}^j$.

Characterizing the conditional mean and variance of returns therefore requires a mapping from the state

variables to valuation ratios. We conjecture, and verify in Section 4.1.1, an affine valuation function conditional on the aggregate state observed at the trading date, $s_t = (Emission_t^{cum}, \pi_t)'$:

$$pd_t^j = a_j + B_j' s_t, \quad B_j = (B_{j,E}, B_{j,\pi})' \quad (13)$$

where π_t is the share of pessimistic investors, introduced and characterized in Section 3.5, and $Emission_t^{cum}$ denotes the cumulative emissions specified in the previous section.

Equations (12) and (13) imply that climate-related state variables affect returns through two distinct mechanisms. The first is cash flows: emissions are a by-product of energy use, so an emission shock moves contemporaneous dividend growth directly through the production technology (3). The second is valuation: because the log price–dividend ratio is affine in the climate state s_t , any revision to cumulative emissions, the distribution of beliefs, and expected climate policy are capitalized into price–dividend ratios. Because price–dividend ratios reflect both expected future dividend growth and future required returns, innovations in price–dividend ratios incorporate both cash-flow news and discount-rate news. Below, we further characterize the cash-flow news and discount-rate news.

Conditional return moments: cash-flow and discount-rate risk

From Equation (12), we can show that return shocks have two sources: news about dividends and news about valuation ratios,

$$\mathbf{r}_{t+1} - E_t[\mathbf{r}_{t+1}] = \underbrace{\Delta d_{t+1} - E_t[\Delta d_{t+1}]}_{\text{cash-flow news}} + \underbrace{\rho (pd_{t+1} - E_t[pd_{t+1}])}_{\text{discount-rate news}} \quad (14)$$

Now, we show how each component is associated with emissions (and by construction climate change). Recall the state vector presented earlier, under the timing introduced in Section 3.5, π_{t+1} is measurable with respect to date- t information; hence the only stochastic element of the state vector s_{t+1} conditional on date t is $Emission_{t+1}^{cum}$, whose innovation is the market-weighted emission innovation. Let $u_{t+1} = (u_{t+1}^B, u_{t+1}^G)'$ collect the (demeaned) asset-level temperature-contribution innovations, with $Var_t(u_{t+1}) = \Sigma_{\Delta T}$.

Cash-flow news. From Equation (3), dividend-growth innovations load on the own-asset element of u_{t+1} with coefficient $(1 - \alpha)$ and on the orthogonal capital shocks with coefficient α : the cash-flow loading on temperature shocks is $L^{CF} \equiv (1 - \alpha)I$.

Discount-rate news. From the affine valuation function (13), valuation-ratio innovations satisfy $pd_{t+1}^j - E_t[pd_{t+1}^j] = B_{j,E} X_t^{M'} u_{t+1}$: discount-rate news is spanned by the same emission shocks that drive temperature, with the rank-one loading $L_t^{DR} \equiv \rho B_E X_t^{M'}$, where $\rho = \text{diag}(\rho_B, \rho_G)$, $B_E = (B_{B,E}, B_{G,E})'$, and X_t^M is the market portfolio.

Collecting the two components, return innovations load on temperature shocks through the matrix

$$L_{r,t} = \underbrace{(1 - \alpha)I}_{\text{cash-flow}} + \underbrace{\rho B_E X_t^{M'}}_{\text{discount-rate}} = L^{CF} + L_t^{DR} \quad (15)$$

and the conditional variance–covariance matrix of returns and the return–temperature covariance decompose accordingly:

$$\Sigma_{r,t} = \underbrace{\alpha^2 \Sigma_K + (1 - \alpha)^2 \Sigma_{\Delta T}}_{\text{cash-flow risk}} + \underbrace{L_t^{DR} \Sigma_{\Delta T} L_t^{DR'}}_{\text{discount-rate risk}} + \underbrace{(1 - \alpha) \left(L_t^{DR} \Sigma_{\Delta T} + \Sigma_{\Delta T} L_t^{DR'} \right)}_{\text{interaction}} = \alpha^2 \Sigma_K + L_{r,t} \Sigma_{\Delta T} L_{r,t}' \quad (16)$$

$$\Sigma_{rT,t} \equiv \text{Cov}_t(\mathbf{r}_{t+1}, u_{t+1}') = \underbrace{(1 - \alpha) \Sigma_{\Delta T}}_{\text{cash-flow exposure}} + \underbrace{L_t^{DR} \Sigma_{\Delta T}}_{\text{discount-rate exposure}} = L_{r,t} \Sigma_{\Delta T} \quad (17)$$

The two components enter with opposite signs. A positive emission shock raises contemporaneous dividends (cash-flow loading $1 - \alpha > 0$) but adds to cumulative emissions and thereby depresses next-period valuations (discount-rate loading $\rho_j B_{j,E} < 0$, since larger inherited cumulative emissions lower valuations, and more so for the brown asset: $B_{B,E} < B_{G,E}$; see Section 4.1.1).² Discount-rate risk thus dampens the total climate exposure of returns relative to its cash-flow component: part of the climate risk that brown assets carry is capitalized into their relative valuations rather than borne through cash flows.

3.3.1 Solution to Investors Optimization

To obtain closed-form expressions for equilibrium expected returns, we assume that investors have constant relative risk aversion (CRRA) preferences and that the underlying shocks are conditionally Gaussian. Combined with the log-linear return approximation in Equation (12) and the affine specification of the log price–dividend

²The statement that carries the results below is the *relative* ordering $B_{B,E} < B_{G,E}$: larger cumulative emissions depress brown valuations relative to green valuations, because expected policy tightening restricts high-emission-intensity production the most and because brown carries greater climate exposure. The level sign of $B_{G,E}$ could be ambiguous in general.

ratio in Equation (13), these assumptions imply that returns are conditionally lognormal. The model therefore falls within the standard CRRA–lognormal framework, which yields tractable closed-form approximations for portfolio demand and equilibrium expected returns. Under these assumptions and after some algebraic manipulation shown in the Appendix, the optimization problem of a type- i investor is:

$$\underset{X_t^{i,G}, X_t^{i,B}}{\text{Max}} E_t^i \left[\frac{\left(\frac{W_{t+1}^i}{1 + \delta \Delta T_t^\eta} \right)^{1-\gamma}}{1-\gamma} \right] \quad (18)$$

$$\text{s.t. } W_{t+1}^i = W_t^i (X'(\mathbf{R}_{t+1} - R_f \mathbf{1}) + R_f)$$

Before solving the optimization problem, we introduce some notations. First, let $\mu_{\Delta T}^i = E_t^i [\log \Delta \mathbf{T}_t]$ and $\Delta T_t^j = v(\text{Emission}_{t+1}^j + \text{Emission}_t^{\text{cum}})$ denote asset j 's contribution to climate change. Before moving forward with finding the solution, we introduce an assumption that helps with easing the exposition.

Assumption. *We assume that both types of investors are fully invested in the risky assets, meaning that $X_t^{i,G} + X_t^{i,B} = 1$. This assumption merely simplifies the exposition by avoiding the need for a Lagrange multiplier in the derivations. In the Online Appendix, we show that our findings remain qualitatively unchanged when we relax this assumption.*

A convenient way to approach the CRRA–lognormal framework is to reduce it to an approximate quadratic form (for reference, see Chapter 3 of Campbell and Viceira (2002)). Since our problem deviates from a standard CRRA–lognormal problem along two dimensions, namely adding a climate damage function and introducing different types with heterogeneous beliefs, two steps are added to the conventional reduction, each relating to one of these new dimensions. We therefore reduce problem (18) to a quadratic problem in the following three steps, detailed in Appendix A.1: (i) the log portfolio return is approximated by $r_{p,t+1}^i \approx r_{f,t} + X_t^{i'}(\mathbf{r}_{t+1} - r_{f,t} \mathbf{1}) + \frac{1}{2} X_t^{i'} \sigma_{r,t}^2 - \frac{1}{2} X_t^{i'} \Sigma_{r,t} X_t^i$; (ii) log damages are linearized around the average temperature, $\log D(\Delta T_t) \approx \bar{\ell} - b(\log \Delta T_t - \overline{\log \Delta T})$, where

$$b = \frac{\eta \delta \overline{\Delta T}^\eta}{1 + \delta \overline{\Delta T}^\eta} \in (0, \eta) \quad (19)$$

is the elasticity of the damage factor with respect to temperature, increasing in the damage convexity η ; and (iii) each investor evaluates damages at the aggregate portfolio weights $X_t = \pi_t X_t^P + (1 - \pi_t) X_t^O$, so that each representative investor internalizes the effect of its own portfolio on aggregate emissions while taking the other

type's portfolio as given. Applying the moment-generating function of the resulting jointly normal variables, the optimization problem of type i reduces to:

$$\underset{X_t^{i,G}, X_t^{i,B}}{Max} \left\{ X_t^{i'} \left(E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{r,t}^2 - b \pi_t^i \mu_{\Delta T}^i \right) - \frac{\gamma}{2} X_t^{i'} \Sigma_{Total,t}^i X_t^i - X_t^{i'} \mathcal{K}_t^i X_t^{i-} \right\} \quad (20)$$

$$where \begin{cases} \gamma \Sigma_{Total,t}^i = \gamma \Sigma_{r,t} + (1 - \gamma) b \pi_t^i (\Sigma_{rT,t} + \Sigma'_{rT,t}) - (1 - \gamma) b^2 (\pi_t^i)^2 \Sigma_{\Delta T} \\ \mathcal{K}_t^i = (1 - \gamma) b \pi_t^{i-} (\Sigma_{rT,t} - b \pi_t^i \Sigma_{\Delta T}) \end{cases}$$

Here $\Sigma_{Total,t}^i$ is the perceived variance–covariance matrix under belief dispersion about climate damages, and the bilinear matrix $\mathcal{K}_t^i$ governs the interaction with the other type's portfolio. $E_t^i[\mathbf{r}_{t+1}]$ denote the expected log return under type i 's belief.

Next, we proceed to solve for the optimal portfolio.

Optimal Portfolio

Solving the optimization problem (20) yields the following optimal portfolio for type i :

$$X_t^i = \underbrace{\frac{1}{\gamma} (\Sigma_{Total,t}^i)^{-1} \left(E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{r,t}^2 - b \pi_t^i \mu_{\Delta T}^i \right)}_{Component\ 1} - \underbrace{\frac{1}{\gamma} (\Sigma_{Total,t}^i)^{-1} \mathcal{K}_t^i X_t^{i-}}_{Component\ 2} \quad (21)$$

where $X_t^i = (X_t^{i,B}, X_t^{i,G})'$ denotes the optimal portfolio of type- i .

The optimal portfolio consists of two components. *Component 1* is the standard mean–variance trade-off adjusted for climate damages. It maximizes a climate-damages adjusted Sharpe ratio. *Component 2* captures the investor's portfolio adjustment in response to the other type's portfolio, and its direction depends on risk aversion through the sign of $\mathcal{K}_t$. For $\gamma < 1$ (and, at high convexity, for $\gamma > 1$, where the damage-variance force dominates), $\mathcal{K}_t$ is positive and the adjustment is *offsetting*: type i views the other type's portfolio as financially suboptimal given climate damages and partially counteracts it (e.g., a pessimist who believes an optimist underinvests in green assets shifts toward green). For $\gamma > 1$ at moderate convexity, $\mathcal{K}_t$ is negative and the adjustment is *reinforcing*: when the other type loads on emissions, aggregate damage risk rises, and an

investor for whom temperature-correlated payoffs are hedges finds those assets more attractive. At log utility ($\gamma = 1$) the interaction vanishes.

Equation (21) characterizes each investor type's optimal portfolio as a *best response* to the portfolio chosen by the other type. In particular, through *Component 2*, type i 's demand depends directly on the holdings of type i^- . The two portfolio choices must therefore be determined jointly as the solution to a system of coupled best responses. The following proposition provides the closed-form interior solution. Let $a_t^i \equiv E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2}\sigma_{\mathbf{r},t}^2 - b\pi_t^i \mu_{\Delta T}^i$ denote type i 's damage-adjusted expected excess return, and define

$$M_t^i \equiv \frac{1}{\gamma} (\Sigma_{Total,t}^i)^{-1} \mathcal{K}_t^i, \quad e_{BMG} \equiv \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \quad e_G \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

Proposition 1 (Optimal portfolios) Suppose $\Sigma_{Total,t}^i$ is positive definite and the spectral radius of M_t^i satisfies $\rho(M_t^O M_t^P) < 1$. Then, the coupled best responses (21) have the unique interior solution

$$\begin{aligned} X_t^P &= (I - M_t^P M_t^O)^{-1} \left[\frac{1}{\gamma} (\Sigma_{Total,t}^P)^{-1} a_t^P - \frac{1}{\gamma} M_t^P (\Sigma_{Total,t}^O)^{-1} a_t^O \right] \\ X_t^O &= (I - M_t^O M_t^P)^{-1} \left[\frac{1}{\gamma} (\Sigma_{Total,t}^O)^{-1} a_t^O - \frac{1}{\gamma} M_t^O (\Sigma_{Total,t}^P)^{-1} a_t^P \right] \end{aligned} \tag{22}$$

Proof: See Appendix A.1.

The condition $\rho(M_t^O M_t^P) < 1$ has a clear economic meaning: it prevents the two types' responses to each other's portfolios (*Component 2*) from becoming explosively strong—each round of mutual adjustment is damped rather than amplified.

As a benchmark, if there are no climate damages (i.e., $D(\Delta T_t) = 1$), then $b = 0$, and the optimal portfolio (21) collapses to the standard mean-variance portfolio:

$$X_t^i = \frac{1}{\gamma} \Sigma_{r,t}^{-1} \left(E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{\mathbf{r},t}^2 \right)$$

3.4 Belief Dispersion

Because our framework is Gaussian, we represent disagreement about climate damages using a convenient “tilting” device: each investor type evaluates expectations under a subjective probability measure that deviates from an objective measure for the climate process. Intuitively, pessimists overweight high-temperature states, while optimists overweight low-temperature states.

Following David (2008), we take the objective measure as the reference and express each type’s beliefs about climate change relative to this benchmark using the Radon–Nikodym derivative. Formally, type i ’s expectation of future temperature change can be written as

$$E_t^i [\Delta T_t] = E_t [\xi^i \Delta T_t] \quad (23)$$

where ξ^i is the Radon–Nikodym derivative that maps the objective probability measure into type i ’s subjective measure. $E_t [\cdot]$ denotes expectations under the objective probability measure. In a Gaussian setting, ξ^i takes the exponential form

$$\xi^i = e^{\theta_i Z_{\Delta T} - \frac{\theta_i^2}{2}} \quad (24)$$

where $Z_{\Delta T}$ denotes the standardized temperature shock under the objective measure. The parameter θ_i governs the direction and strength of the tilt of probabilities: if $\theta_i > 0$, type i places more probability weight on higher-temperature realizations (climate pessimism); if $\theta_i < 0$, type i places more weight on lower-temperature realizations (climate optimism).

A useful implication is that disagreement manifests as a mean shift in perceived future temperature change. We show in Online Appendix that

$$\mu_{\Delta T}^P = \mu_{\Delta T} + \sigma_{\Delta T} \theta_P > \mu_{\Delta T} \quad (25)$$

$$\mu_{\Delta T}^O = \mu_{\Delta T} + \sigma_{\Delta T} \theta_O < \mu_{\Delta T} \quad (26)$$

$$\text{where } \theta_O < 0 < \theta_P$$

Here $\mu_{\Delta T}$ and $\sigma_{\Delta T}$ are the objective mean and standard deviation of temperature change. Equations (25) and (26) show that pessimistic investors perceive a higher expected temperature change than the objective benchmark (i.e., $\mu_{\Delta T}^P > \mu_{\Delta T}$), while optimistic investors perceive a lower one (i.e., $\mu_{\Delta T}^O < \mu_{\Delta T}$).

Climate beliefs and return beliefs

With an endogenous price–dividend ratio, future valuation ratios and dividends depend on climate states—cumulative emissions, beliefs, and policy—so the probability tilt θ_i that shifts perceived temperature generally also shifts perceived returns. Specifically, each type applies its tilt to the temperature innovation wherever it appears. Because return innovations load on temperature shocks through $L_{r,t}$ in (15), type i ’s return beliefs satisfy

$$E_t^i[\mathbf{r}_{t+1}] = E_t[\mathbf{r}_{t+1}] + \theta_i L_{r,t} \sigma_{\Delta T} \quad (27)$$

The return-belief wedge inherits the two opposing components of $L_{r,t}$. Through the cash-flow component $(1 - \alpha)I$, a pessimist—who overweights high emission realizations—perceives *higher* contemporaneous energy-driven dividends; through the valuation component $\rho B_E X_t^{M'}$, the same pessimist perceives larger future cumulative emissions and hence *lower* future valuations, especially for brown assets. The net direction of the wedge is therefore an equilibrium question rather than an assumption. Investors continue to share common beliefs about productivity and capital shocks; disagreement concerns only climate-related innovations, which affect returns through both cash flows and valuations.

Considering the belief characterization and combining it with the optimal portfolio structure derived in the previous section, the difference between the two types’ portfolios (i.e., the demand gap) is

$$X_t^O - X_t^P = \frac{e'_{BMG} \mathcal{N}}{e'_{BMG} \mathcal{D} e_{BMG}} e_{BMG} \quad (28)$$

$$\text{where } \mathcal{N} = b [\pi_t \mu_{\Delta T}^P - (1 - \pi_t) \mu_{\Delta T}^O] - (\theta_P - \theta_O) L_{r,t} \sigma_{\Delta T} + [\gamma \Sigma_{\text{Total},t}^P + \mathcal{K}_t^P - \gamma \Sigma_{\text{Total},t}^O - \mathcal{K}_t^O] X_t^M$$

$$\text{and } \mathcal{D} = [\gamma(1 - \pi_t) \Sigma_{\text{Total},t}^P + \gamma \pi_t \Sigma_{\text{Total},t}^O - \pi_t \mathcal{K}_t^P - (1 - \pi_t) \mathcal{K}_t^O]$$

The demand gap (28) reflects how the same climate beliefs affect both perceived damages and perceived returns. The first term in the numerator captures the damage channel: a pessimist who overweights high-temperature realizations perceives greater damages and therefore has a stronger incentive to reduce brown exposure, other

things equal. Population shares also matter because each type's portfolio enters aggregate temperature with its respective weight. A larger group can induce a greater reduction in aggregate temperature through the same shift toward green assets, so the damage channel depends on both perceived damages and the group's influence on aggregate emissions.

The second term captures the return-belief wedge described above. When the cash-flow component raises pessimists' perceived brown returns relative to green returns, it offsets their incentive to reduce brown exposure. When the valuation component instead lowers their perceived relative brown returns, it reinforces that incentive. The net effect depends on the brown–green loading in $L_{r,t}\sigma_{\Delta T}$.

The third term captures differences in how portfolio risk and the interaction between the two types' holdings affect their desired allocations. Although both types face the same underlying shocks, their different population shares imply that a given portfolio adjustment changes aggregate temperature by different amounts. This changes the balance between financial returns and climate damages, including the covariance between them. The demand gap therefore also depends on the brown exposure in the market portfolio, rather than on the belief difference $\theta_P - \theta_O$ alone.

Having characterized investors' beliefs, the next step is to determine the share of each generation that adopts a given belief. We address this in the subsequent section.

3.5 Endogenous Shares of Optimists and Pessimists

In this section, we describe how the shares of climate optimists and pessimists endogenously evolve over time. To this end, we define π_t (resp. $1 - \pi_t$) as the share of pessimist (resp. optimist) investors, and endogenize π_t within the model. Specifically, we follow the literature on adaptive beliefs and allow π_t to evolve based on past relative performance (see e.g., Brock and Hommes (1998)). In particular, the fraction of pessimists and optimists in generation $t + 1$ depends on the realized relative financial performance of pessimists versus optimists in the most recently completed generation, $t - 1$. Intuitively, if one type performs better financially, a larger fraction of investors adopts that type's belief in the subsequent generation. The one-generation lag is not merely a convention: it is what makes the belief-adoption dynamics well-defined once prices are endogenous, because the performance of generation t is not realized until the date- $t + 1$ price is set, and that price depends on the composition of the entering cohort. Sorting the entering cohort on the last completed generation's

performance removes this simultaneity.

We measure past relative performance using the difference in realized utility between pessimist and optimist investors. To do so, we define $\hat{W}_{t+1}^i = W_t^i (X'(\mathbf{R}_{t+1} - R_f \mathbf{1}) + R_f) \times D(\Delta T_t)$. The difference in realized utility between optimists and pessimists is then given by:

$$\Delta U_t = U(\hat{W}_{t+1}^P) - U(\hat{W}_{t+1}^O) \quad (29)$$

We consider the following logit belief adoption dynamics, in which the entering cohort at $t + 1$ sorts on the performance of the last completed generation:

$$\pi_{t+1} = \frac{\pi_t e^{\beta \Delta U_{t-1}}}{\pi_t e^{\beta \Delta U_{t-1}} + (1 - \pi_t)} \quad (30)$$

where $\beta \geq 0$ measures how strongly investors respond to relative performance when choosing their type. When $\beta = 0$, we have $\pi_{t+1} = \pi_t$, implying that investors disregard past performance. As $\beta \rightarrow \infty$, all new investors adopt the type with higher realized utility. Under relatively smooth fluctuations in utility (i.e., small ΔU_{t-1}), we can linearly approximate (30) as below:

$$\pi_{t+1} \approx \pi_t + \beta \pi_t (1 - \pi_t) \Delta U_{t-1} \quad (31)$$

Moreover, ΔU_{t-1} can be approximated by:

$$\Delta U_{t-1} \approx U'(\bar{W}_t) (\hat{W}_t^P - \hat{W}_t^O) \quad (32)$$

where $\bar{W}_t = \frac{\hat{W}_t^O + \hat{W}_t^P}{2}$.

To further clarify how belief adoption dynamics operate, we describe the economy at time $t + 1$.

First, the new cohort enters already sorted: its composition π_{t+1} was determined at date t by the logit rule (30), applied to the realized relative performance of generation $t - 1$ —exploiting the no-information-asymmetry assumption, which lets entering investors observe the previous generations' portfolios and realized payoffs.

Second, the market trade takes place: exiting generation t investors are forced sellers (they need cash to consume), and entering generation $t + 1$ investors, sorted by type, are the only available buyers. Cash flows from

generation $t + 1$ to generation t , and the float of brown and green claims flows the other way; the price–dividend ratios PD_{t+1} adjust to make generation $t + 1$'s type-weighted demand absorb the entire supply generation t is unloading.

Third, the trade pins down prices P_{t+1} and hence generation t 's realized returns r_{t+1} ; together with generation t 's chosen portfolios X_t^P, X_t^O (set at t under their respective beliefs), this determines the realized terminal wealths $\hat{W}_{t+1}^P, \hat{W}_{t+1}^O$ and therefore makes $\Delta U_t \equiv U(\hat{W}_{t+1}^P) - U(\hat{W}_{t+1}^O)$ computable. By (30), ΔU_t sets the composition of the *next* entering cohort, π_{t+2} .

This dynamic is depicted in Figure 3, which visualizes the interaction between generation t and $t + 1$. Under this timing, generation $t + 1$'s demand affects prices at $t + 1$ and hence r_{t+1} , and r_{t+1} feeds into π_{t+2} , which is predetermined when generation $t + 1$ trades.³




Figure 3. Timeline of overlapping generations and belief adoption

³One within-date simultaneity worth noting: the date- t valuation pd_t capitalizes $E_t[pd_{t+1}]$, which contains the date- t -measurable π_{t+1} , while π_{t+1} depends on r_t and hence on pd_t . This is a standard within-period fixed point in equations—no agent requires information that does not yet exist—and it has a unique solution under the contraction condition $\Lambda_t \rho_{\max} |B\pi|_{\max} < 1$, which bounds the product of the belief sensitivity and the price sensitivity to beliefs.

Combining Equations (31) and (32), we obtain the following lemma.

Lemma 1. *In an economy with otherwise identical pessimistic and optimistic investors and a short-sale constraint, a higher realized BMG return reduces the share of pessimists in the next entering cohort. Formally,*

$$\pi_{t+1} \approx \pi_t - \Lambda_t r_t^{BMG}, \quad (33)$$

$$\text{where } \Lambda_t \equiv \beta \pi_t (1 - \pi_t) U'(\bar{W}_t) W_{t-1} \bar{D}_{t-1} \Delta X_{t-1} \geq 0$$

Λ_t can be regarded as *belief sensitivity coefficient*, $\bar{D}_{t-1}$ denotes the realized damage factor applied to generation $t - 1$'s terminal wealth, and $r_t^{BMG} = r_t^B - r_t^G$ denotes the realized log BMG return from $t - 1$ to t . Every quantity in Λ_t is dated t or earlier, so the rule is predetermined with respect to date- $t + 1$ trading. Also, $\Delta X_{t-1}^B = |\Delta X_{t-1}^G| = \Delta X_{t-1}$ which denotes the difference in portfolio holdings between optimists and pessimists in generation $t - 1$, with $\Delta X_{t-1}^B > 0$ since optimists hold a larger share of brown assets than otherwise identical pessimists.^{4 5}

4 Equilibrium

In this section, we characterize equilibrium expected returns under two scenarios: (i) the absence of climate policy, and (ii) the presence of climate policy in the form of an emissions cap. We also analyze the properties of the brown-minus-green (BMG) portfolio in equilibrium. These results help reconcile the disconnect between existing theoretical predictions and empirical evidence, supporting our central argument that belief dispersion regarding climate damages constitutes a key missing state variable.

⁴Consider that investors are otherwise identical (i.e., they have the same initial endowment and preferences) and both types are fully invested in risky assets, i.e., $X_t^{O,B} + X_t^{O,G} = X_t^{P,B} + X_t^{P,G} = 1$. Subtracting yields $\Delta X_t^B + \Delta X_t^G = 0$. Since optimists hold a larger share of brown assets than otherwise identical pessimists, it follows that $\Delta X_t^B = |\Delta X_t^G| > 0$.

⁵The reason a pessimist holds a smaller share of the brown asset in her portfolio than an otherwise identical optimist is that, for a given allocation between brown and green assets, the pessimist derives lower utility due to her higher perceived climate damages associated with brown investment. Consequently, she is willing to reduce her exposure to the brown asset, even at a discount. In contrast, the optimist—who perceives lower damages—finds the discounted brown asset attractive and is willing to increase his holdings. This trade benefits both parties: the pessimist improves her utility by reducing exposure to perceived high-damage assets, while the optimist benefits from acquiring underpriced brown assets. As a result, in equilibrium, pessimists hold a smaller share of brown assets than otherwise identical optimists. The demand gap (28) in Proposition 1 gives the closed form of this holding difference.

4.1 Equilibrium under No Climate Policy

We first derive equilibrium outcomes when no emissions cap is imposed. Market clearing requires that the aggregate demand of pessimists and optimists equals the market weight of each risky asset:

$$\pi_t X_t^{P,B} + (1 - \pi_t) X_t^{O,B} = X_t^B \quad (34)$$

$$\pi_t X_t^{P,G} + (1 - \pi_t) X_t^{O,G} = X_t^G \quad (35)$$

where X_t^B and X_t^G denote the market portfolio weights of the brown and green assets in generation t , respectively, and π_t is the fraction of pessimists that was endogenously derived in the previous section. Market-clearing condition (34) implies that the aggregate holdings of pessimists and optimists in the brown asset equal its supply. Similarly, condition (35) implies that the aggregate holdings of both types in the green asset equal its supply.

Combining (34)–(35) with belief dispersion structure and the optimal portfolio, we derive closed-form solutions for the assets' equilibrium expected returns.

Equilibrium Expected Returns

Define $\mu_{r,t} \equiv E_t[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{\mathbf{r},t}^2$ as the equilibrium expected return (The addition of the one-half variance term ($\frac{1}{2} \sigma_{\mathbf{r},t}^2$) is necessary to convert log-return to simple return, which is the ultimate concern of the representative investor) and applying the market clearing condition yields:

$$\begin{aligned} \mu_{r,t} = & \gamma \left[\pi_t \Sigma_{\text{Total},t}^P X_t^P + (1 - \pi_t) \Sigma_{\text{Total},t}^O X_t^O \right] \\ & + b \left[\pi_t^2 \mu_{\Delta T}^P + (1 - \pi_t)^2 \mu_{\Delta T}^O \right] \\ & + \pi_t \mathcal{K}_t^P X_t^O + (1 - \pi_t) \mathcal{K}_t^O X_t^P \\ & - [\pi_t \theta_P + (1 - \pi_t) \theta_O] L_{r,t} \sigma_{\Delta T}. \end{aligned} \quad (36)$$

Then, expanding and simplifying (36), we show in the Appendix that equilibrium expected excess returns

under the objective measure has the following structure:

$$\begin{aligned}
\mu_{r,t} = & \underbrace{\gamma \alpha^2 \Sigma_K X_t^M}_{\text{financial risk premium}} + \underbrace{\gamma L_{r,t} \Sigma_{\Delta T} L'_{r,t} X_t^M}_{\text{climate risk premium}} \\
& + \underbrace{b \left(\pi_t^2 \mu_{\Delta T}^P + (1 - \pi_t)^2 \mu_{\Delta T}^O \right)}_{\text{speculation on climate beliefs}} \\
& + \underbrace{(1 - \gamma) b \left\{ \Sigma_{rT,t} X_t^M + \Sigma'_{rT,t} \left[\pi_t^2 X_t^P + (1 - \pi_t)^2 X_t^O \right] - b \left[\pi_t^2 + (1 - \pi_t)^2 \right] \Sigma_{\Delta T} X_t^M \right\}}_{\text{Distorted risk sharing}} \\
& - \underbrace{(\pi_t \theta_P + (1 - \pi_t) \theta_O) L_{r,t} \sigma_{\Delta T}}_{\text{Belief-driven valuation}}
\end{aligned} \tag{37}$$

Equation (37) is already a giant equation, so to explain it, we simplify some terminologies. In what follows, we write $\tilde{\theta}_t \equiv \pi_t \theta_P + (1 - \pi_t) \theta_O$ for the average probability tilt, and refer to the term associated with the distorted-risk-sharing as *Hedging*_t.

Given these notational simplifications, Equation (37) admits a clear economic decomposition. The first two terms are conventional risk premia: they compensate investors for exposure to financial and climate risk, scale with risk aversion, and together form a standard two-factor benchmark. The remaining three terms arise because investors internalize climate damages under their own subjective beliefs about temperature change. They capture speculation on climate beliefs, distorted risk sharing, and a belief-driven valuation wedge. Their magnitudes depend on the distribution of beliefs and on each type's population share.

First, $\gamma \alpha^2 \Sigma_K X_t^M$ is the *financial risk premium*: compensation for exposure to the capital-productivity shocks Σ_K . Because capital and energy shocks are orthogonal by assumption, this term is orthogonal to emission shocks—it is the pure financial factor, unaffected by climate states.

Second, $\gamma L_{r,t} \Sigma_{\Delta T} L'_{r,t} X_t^M$ is the *climate risk premium*: compensation for the component of ordinary return risk generated by climate-related innovations u_{t+1} , through their effects on dividends and future valuations.

Third, $b \left(\pi_t^2 \mu_{\Delta T}^P + (1 - \pi_t)^2 \mu_{\Delta T}^O \right)$ captures *speculation on climate beliefs*. Both types internalize the damages their financed emissions cause, but each does so under its own subjective measure, so what is priced is

the share-weighted average perceived warming, scaled by b , the sensitivity of wealth to temperature. Because the two types evaluate the same assets under different measures, they willingly take opposite sides of climate-related risk (Harrison and Kreps, 1978): the optimist buys brown exposure from the pessimist at a price each regards as favorable. This term is the counterpart of $Speculation_t^j$ in Proposition 2. It carries no factor of $(1 - \gamma)$, so it prices climate at every level of risk aversion.

Fourth, $Hedging_t$ captures *distorted risk sharing*. Its first term, $\Sigma_{rT,t} X_t^M$, reflects investors adjusting their portfolios in response to the covariance between asset returns and aggregate temperature. The second term, $\Sigma'_{rT,t} [\pi_t^2 X_t^P + (1 - \pi_t)^2 X_t^O]$, captures how each representative type's investment changes the covariance between its portfolio returns and aggregate temperature. The third term, $-b [\pi_t^2 + (1 - \pi_t)^2] \Sigma_{\Delta T} X_t^M$, captures how investment changes aggregate temperature variance. Together, these effects shape portfolio choices because investors care about damage-adjusted wealth. Their signs depend on risk aversion, damage sensitivity, population shares, and the covariance structure.

Fifth, $-\tilde{\theta}_t L_{r,t} \sigma_{\Delta T}$ is the *belief-driven valuation* term. With endogenous price–dividend ratios, climate beliefs spill over into return beliefs through (27), so an average tilt $\tilde{\theta}_t$ shifts perceived return prospects by $\tilde{\theta}_t L_{r,t} \sigma_{\Delta T}$. That perception is expressed as demand, demand moves prices, and objective-measure expected returns must adjust in the opposite direction for the market to clear. Unlike the first two terms, this is not compensation for bearing risk. Instead, it is a wedge between the measure under which assets are priced and the measure under which returns are realized.

Factor Model Representation

Equation (37) can also be written in a conventional factor-pricing form. We establish the following proposition in the Appendix.

Proposition 2 *Equilibrium expected returns are characterized by a two-factor model – financial and climate factors—augmented by belief-driven pricing terms:*

$$\mu_{r,t}^j = \beta_{F,t}^j \lambda_{F,t} + \beta_{C,t}^j \lambda_{C,t} + b \times Speculation_t^j + Hedging_t^j + Valuation_t^j \quad (38)$$

$$where \ Speculation_t^j = [\pi_t^2 \mu_{\Delta T}^P + (1 - \pi_t)^2 \mu_{\Delta T}^O]_j \quad ; \quad Valuation_t^j = -\tilde{\theta}_t [L_{r,t} \sigma_{\Delta T}]_j$$

$$Hedging_t^j = \left[(1 - \gamma)b\{\Sigma_{rT,t}X_t^M + \Sigma'_{rT,t} \left[\pi_t^2 X_t^P + (1 - \pi_t)^2 X_t^O \right] - b \left[\pi_t^2 + (1 - \pi_t)^2 \right] \Sigma_{\Delta T} X_t^M \} \right]_j$$

Here, $\beta_{F,t}^j$ and $\beta_{C,t}^j$ denote the exposures of asset j to the financial and climate risk factors, respectively, while $\lambda_{F,t}$ $\lambda_{C,t}$ denote the corresponding prices of risk for these two systematic factors. The two factors are defined on the two blocks of the risk compensation in (37). The *financial factor* is built on the capital-shock covariance matrix Σ_K . The *climate factor* is built on the climate-state risk matrix

$$\Theta_t \equiv \gamma L_{r,t} \Sigma_{\Delta T} L'_{r,t} \quad (39)$$

which collects everything that rewards exposure to emissions shocks and as characterized in Equations (16) and (17), captures the cash-flow news, discount-rate news, and their interaction. The last three terms capture (i) disagreement-driven speculation, (ii) the equilibrium hedging demand induced by belief-driven distortions in optimal risk sharing, and (iii) the return-belief wedge. In the special case of no belief dispersion, the speculation and wedge terms collapse to their common-belief counterparts and the model reduces to the standard two-factor benchmark.

4.1.1 Determination of the Price–Dividend Ratio

Combining the definition of $\mu_{r,t}$ with the log-linear return identity (12) yields, asset by asset,

$$pd_t^j = k_j - r_f + E_t[\Delta d_{t+1}^j] + \rho_j E_t[pd_{t+1}^j] - \mu_{r,t}^j + \frac{1}{2} \sigma_{j,t}^2 \quad (40)$$

Iterating (40) forward under the transversality condition $\lim_{s \rightarrow \infty} \rho_j^s E_t[pd_{t+s}^j] = 0$ gives the present-value representation

$$pd_t^j = \sum_{s \geq 0} \rho_j^s E_t \left[k_j - r_f + \Delta d_{t+1+s}^j - \mu_{r,t+s}^j + \frac{1}{2} \sigma_{j,t+s}^2 \right]$$

Valuations are high when expected dividend growth is high and when future equilibrium risk compensations are low.

The affine conjecture (13) is verified by the method of undetermined coefficients. Let the linearized law of motion of the state be $s_{t+1} = c_t + \Phi s_t + H u_{t+1}$ where Φ is 2×2 matrix, and let $E_t[\Delta d_{t+1}^j] = g_{0j} + g_j' s_t$ and $\mu_{r,t}^j = m_{0j} + m_j' s_t$, where the latter follows from (37) under the linearization with state-independent

conditional second moments. We establish the following proposition in the Appendix.

Proposition 3 *The affine valuation function (13) solves (40) if and only if the loadings satisfy*

$$\begin{aligned} B'_j(I - \rho_j \Phi) &= g'_j - m'_j, \\ a_j(1 - \rho_j) &= k_j - r_f + g_{0j} - m_{0j} + \rho_j B'_j c + \frac{1}{2} \sigma_j^2. \end{aligned} \tag{41}$$

If $0 < \rho_j < 1$ and $\rho_j \rho(\Phi) < 1$, these equations determine the affine coefficients:

$$B'_j = (g'_j - m'_j)(I - \rho_j \Phi)^{-1}, \quad a_j = \frac{k_j - r_f + g_{0j} - m_{0j} + \rho_j B'_j c + \frac{1}{2} \sigma_j^2}{1 - \rho_j},$$

where $\rho(\Phi)$ denotes the spectral radius of the state-transition matrix.

Because the coefficients $(g_j, m_j, \sigma_j^2, \Phi)$ themselves depend on (a_j, B_j) through the second moments (16) and through ρ_j , Equation (41) is a fixed-point system in the loadings; existence follows from continuity on a compact set of candidate loadings, while uniqueness should be verified numerically, since self-fulfilling valuation-risk equilibria cannot be ruled out in general.

The comparative content of (41) is as follows: The cumulative-emissions loading satisfies $B_{j,E} \propto g_{j,E} - m_{j,E}$: higher cumulative emissions lower expected dividend growth (expected future caps tighten, restricting high-emission-intensity production the most) and raise required returns (expected damages and the marginal damage sensitivity b rise with cumulative emissions under convex damages), so $B_{B,E} < B_{G,E}$, with $B_{B,E} < 0$ unambiguously. Because cumulative emissions are highly persistent, these effects are capitalized with the factor $1/(1 - \rho_j)$ rather than discounted away. The belief loading $B_{j,\pi}$ inherits the sign of $-\partial \mu_{r,t}^j / \partial \pi_t$: a larger share of pessimists raises the required brown return and lowers the required green return, so $B_{B,\pi} < 0 < B_{G,\pi}$.

4.1.2 BMG Portfolio

Having derived the structure of equilibrium expected returns, we now characterize the expected return on the brown-minus-green (BMG) portfolio and show how belief dispersion rationalizes both sign changes and time variation in expected returns.

Property 1: *Under the parameter conditions in Appendix B.4, the equilibrium BMG premium can be written*

as:

$$\mu_{r,t}^{BMG} = \mu_{r,t}^{BMG,0} + \mathcal{B}_t^P \theta_P + \mathcal{B}_t^O \theta_O \quad (42)$$

where $\mu_{r,t}^{BMG,0}$, $\mathcal{B}_t^P$, and $\mathcal{B}_t^O$ are determined in the Appendix. Then the sign of BMG portfolio expected return is:

- Positive when $|\theta_O| < \bar{\theta}_t + \chi_t \theta_P$, i.e., when optimists' beliefs do not deviate strongly from pessimists' beliefs;
- Negative when $|\theta_O| > \bar{\theta}_t + \chi_t \theta_P$, i.e., when optimists' beliefs deviate strongly from pessimists' beliefs;
- Zero when $|\theta_O| = \bar{\theta}_t + \chi_t \theta_P$,

where $\bar{\theta}_t \equiv \frac{\mu_{r,t}^{BMG,0}}{\mathcal{B}_t^O}$ and $\chi_t \equiv \frac{\mathcal{B}_t^P}{\mathcal{B}_t^O}$.

Property 1 makes the mechanism explicit: because disagreement alters both speculative demand and the belief-driven hedging demand, the BMG premium need not be stable and can flip sign as belief dispersion changes. Put differently, when disagreement becomes sufficiently strong, the speculative and externality components can dominate the conventional climate-risk compensation embedded in the BMG portfolio.

This observation motivates a natural follow-up question: when would both types—optimists and pessimists—simultaneously discount brown assets and bid up green assets, eliminating the possibility that speculative motives overturn the sign?

To address this question, consider the convexity of climate damages governed by η . Higher η means economic damages rise more than proportionally with temperature, so marginal warming becomes increasingly costly. To develop the property associated with damage convexity, we first present an assumption and justification behind it. We assume that the brown asset has higher sensitivity to climate (or emissions) shocks than the green assets, that is $\sigma_{\Delta T}^B > \sigma_{\Delta T}^G > 0$. The root of this assumption can be found in Equations (4) and (5). To clarify, both sectors emit because they use energy in their production activity. In this sense, when the two sectors face a common activity shock and have comparable production responses, the emissions shock attributed to the brown asset is larger: as shown in (5), the brown asset has a higher emissions intensity, so the same activity shock translates into a larger emissions shock and consequently a larger climate shock. Given this reasonable assumption, we introduce the following property.

Property 2: For $\gamma > 1$, there exists a finite threshold $\bar{\eta}$ such that the BMG portfolio earns a positive expected return whenever $\eta > \bar{\eta}$ (i.e., when the damage is sufficiently convex) where $\bar{\eta}$ is determined in the Appendix.

Property 2 states that when damages are sufficiently convex, marginal warming becomes costly enough that *both* types—not just pessimists—heavily discount the brown asset, and the BMG premium becomes positive regardless of belief dispersion.

4.2 Equilibrium with Climate Policy

We now introduce climate policy into the model in the form of an emissions cap. Specifically, we incorporate the cap constraint as an additional equilibrium condition. The level of the cap is determined by political pressure exerted by both types of investors, where the extent of pressure depends on their respective climate beliefs. We then examine how policy shifts induced by this political pressure affect expected returns. Because political pressure governs the stringency of the future emissions cap, it serves as the model's fundamental source of transition risk. In particular, a tightening cap alters firms' feasible emissions and, consequently, their production possibilities, with first-order implications for brown assets.

4.2.1 Climate Policy

The policymaker may intervene by limiting permissible temperature warming through an aggregate emissions cap. This emissions cap responds to political pressure generated by each type of investor. In particular, we consider a probabilistic voting economy. Therefore, given that π_t denote the share of pessimistic investors and $1 - \pi_t$ the share of optimistic investors, we define political pressure, ϕ_t , as

$$\phi_t = \pi_t \frac{E_t^P [\Delta T_t]}{\overline{\Delta T}} - (1 - \pi_t) \frac{E_t^O [\Delta T_t]}{\overline{\Delta T}} \quad (43)$$

where $\overline{\Delta T}$ is the long-run average of temperature change. Then the cap is determined by:

$$Cap_{t+1} = \overline{Cap} \times e^{-\psi \phi_t} \quad (44)$$

where ψ measures the responsiveness of policy to belief-driven political pressure. When imposed, the cap constrains aggregate emissions:

$$Cap_{t+1} = X_t^G Emission_{t+1}^G + X_t^B Emission_{t+1}^B \quad (45)$$

A lower Cap_t indicates a more stringent climate policy, while a higher cap is more lenient.

4.2.2 Equilibrium Relations

Considering the emissions cap and its relationship with the political pressure, we develop the following proposition, which characterizes how the expected returns of brown and green assets respond to political pressure. Before introducing the Proposition, it is useful to provide intuition on what changes when the level of climate pessimism and consequently the political pressure shift. More pessimists affect the market in three ways: First is the investor composition. If pessimists hold less brown than optimists, replacing optimists with pessimists reduces brown demand. This pushes required BMG compensation upward. Second is through emissions cap. More pessimists generate stronger political pressure and a tighter cap. Through the cap Equation (45), this reduces the brown market weight investors must collectively hold. Having less brown exposure to absorb pushes required compensation downward. Third is through changes in portfolio incentives. Changing population shares changes each representative type's marginal influence on aggregate temperature. Changing the market portfolio also changes return loadings. These adjustments can strengthen or weaken the reduction in brown demand.

Bearing these three channels in mind, we develop the proposition that summarizes the relationship between BMG return and political pressure. To avoid excessively large expressions, we define a few variables that are used in expressing Proposition 4. Let $\mathcal{A}_t^i$ and $\mathcal{J}_t^i$ denote type i 's portfolio curvature in the brown–green direction and how the other type's reallocation from green to brown changes type i 's required BMG compensation, respectively:

$$\mathcal{A}_t^i \equiv \gamma e'_{BMG} \Sigma_{\text{Total},t}^i e_{BMG} \quad \mathcal{J}_t^i \equiv e_{BMG}' \mathcal{K}_t^i e_{BMG}$$

Also, G_t denote the change in type i 's required BMG compensation due to changing coefficients, evaluated at

its current holdings:

$$G_t^i \equiv e'_{BMG} \left[\gamma \frac{d\Sigma_{\text{Total},t}^i}{d\pi_t} X_t^i + \frac{d\mathcal{K}_t^i}{d\pi_t} X_t^{i-} + b \frac{d\pi_t^i}{d\pi_t} \mu_{\Delta T}^i - \theta_i \frac{dL_{r,t}}{d\pi_t} \sigma_{\Delta T} \right], \quad (46)$$

and S_t denote the slope of BMG compensation with respect to the brown market weight when portfolio coefficients are held fixed:

$$\mathcal{S}_t \equiv \frac{\mathcal{A}_t^P \mathcal{A}_t^O - \mathcal{J}_t^P \mathcal{J}_t^O}{e'_{BMG} \mathcal{D} e_{BMG}} > 0 \quad (47)$$

where $\mathcal{D}$ was defined when we introduced demand gap in Equation (28). Lastly, let $\mathcal{R}_t$ denote the adjustment in required BMG compensation arising from changes in each type's portfolio incentives as the pessimistic share increases. Changing the mix of investors affects aggregate brown demand even if each type maintains its initial portfolio; this composition effect is captured separately by the holding gap $\Delta X_t = X_t^{O,B} - X_t^{P,B}$. However, population shares also change each representative type's marginal influence on aggregate temperature, thereby altering its incentive to internalize climate damages. In addition, the associated tightening of the emissions cap changes the market portfolio and return loadings, affecting portfolio risk and perceived returns. Together, these changes lead both types to revise their desired portfolios. $\mathcal{R}_t$ captures the resulting adjustment in required BMG compensation after accounting for their joint portfolio responses. Given all that said, we have:

$$\mathcal{R}_t \equiv \frac{[\pi_t \mathcal{A}_t^O - (1 - \pi_t) \mathcal{J}_t^O] G_t^P + [(1 - \pi_t) \mathcal{A}_t^P - \pi_t \mathcal{J}_t^P] G_t^O}{e'_{BMG} \mathcal{D} e_{BMG}}. \quad (48)$$

Proposition 4 *Under the assumptions in Appendix B.3, an increase in political pressure generated by a larger pessimistic share changes the BMG premium according to:*

$$\frac{d\mu_{r,t}^{BMG}}{d\phi_t} = \mathcal{S}_t \left[\frac{\Delta X_t}{d\phi_t/d\pi_t} - \frac{\psi Cap_{t+1}}{Emission_{t+1}^B - Emission_{t+1}^G} \right] + \frac{\mathcal{R}_t}{d\phi_t/d\pi_t} \quad (49)$$

where $d\phi_t/d\pi_t = \frac{E_t^P[\Delta T] + E_t^O[\Delta T]}{\Delta T} > 0$. Because $\mathcal{S}_t > 0$, the premium increases if and only if

$$\underbrace{\Delta X_t + \frac{\mathcal{R}_t}{\mathcal{S}_t}}_{\text{Term 1}} > \underbrace{\frac{\psi Cap_{t+1}}{Emission_{t+1}^B - Emission_{t+1}^G} \frac{d\phi_t}{d\pi_t}}_{\text{Term 2}}$$

where Term 1 captures how much investors want to reduce brown holdings as pessimism increases, and Term 2 captures how much the tighter cap reduces the brown weight in the market portfolio.

Proof: See the Appendix.

Proposition 4 characterizes two opposing effects: First, investors may want to hold less brown. More investors adopt pessimistic portfolios, and both types may revise their holdings as their influence on aggregate temperature and perceived returns changes. This reduction in demand pushes required BMG compensation up.

Second, greater pessimism strengthens political pressure and tightens the emissions cap. Through your model's cap equation, this reduces the brown market weight investors must collectively hold. Having less brown exposure to absorb pushes required BMG compensation down.

When the former dominates, the BMG return rises in response to political pressure, and the latter dominates, it reverses.

Next, we examine the return predictability generated by the belief-policy feedback. With endogenous prices, a distinction becomes central: news about future beliefs and policy is capitalized into current valuations, whereas predictability of future returns is governed solely by how the belief composition moves future required compensations. Using Lemma 1, we introduce Proposition 5, which highlights how endogenous changes in the shares of optimists and pessimists affect political pressure and, consequently, climate policy. This interaction generates a feedback loop that drives BMG return dynamics.

Proposition 5 *Under the belief-adoption dynamics of Lemma 1, the effect of realized BMG performance on the next generation's required BMG compensation is*

$$\frac{d\mu_{r,t+1}^{BMG}}{dr_t^{BMG}} = -\Lambda_t \left\{ \mathcal{S}_{t+1} \left[\Delta X_{t+1} - \frac{\psi Cap_{t+2}}{Emission_{t+2}^B - Emission_{t+2}^G} \frac{d\phi_{t+1}}{d\pi_{t+1}} \right] + \mathcal{R}_{t+1} \right\} \quad (50)$$

If $\Lambda_t > 0$ and $\mathcal{S}_{t+1} > 0$, we observe reversal when

$$\Delta X_{t+1} + \frac{\mathcal{R}_{t+1}}{\mathcal{S}_{t+1}} > \frac{\psi Cap_{t+2}}{Emission_{t+2}^B - Emission_{t+2}^G} \frac{d\phi_{t+1}}{d\pi_{t+1}}. \quad (51)$$

The reverse inequality produces momentum

Proof: See the Appendix.

The intuition is as follows. Provided optimists previously held more brown assets than pessimists, a higher

realized BMG return improves their relative performance and shifts the next entering cohort toward optimism (Lemma 1). The resulting decline in the pessimistic share weakens political pressure and relaxes the emissions cap. This creates two opposing effects. Greater optimism can increase investors' willingness to hold brown assets, lowering required BMG compensation. However, through the cap equation, a more lenient policy increases the brown market weight that investors must collectively hold, raising the compensation needed to absorb that exposure. Changes in population shares also alter each type's marginal damage internalization, while changes in the market portfolio affect return loadings.

Proposition 5 therefore suggests a conditional reversal. When the increase in desired brown holdings, including these incentive adjustments, exceeds the cap-induced increase in the brown market weight, required BMG compensation falls. Higher past BMG returns then lead to lower expected BMG compensation at the next generation's investment horizon. Momentum arises when the increase in required brown holdings exceeds the adjusted increase in demand. It does not require the valuation-belief wedge alone to dominate: the cap response, damage internalization, and portfolio interactions can all affect the outcome.

4.2.3 Extension: Median Voting Economy

Previously, we considered a probabilistic voting model that determines climate policy. In this section, we analyze the implications of another well-known voting model: median voting. Under median voting the cap is not a smooth weighted average of political pressures; it is instead the policy preferred by the median voter. Let the preferred cap of type $i \in \{O, P\}$ be

$$Cap_{t+1}^i = \overline{Cap} \times e^{-\psi \frac{E_t^i[\Delta T_t]}{\Delta T}} \quad (52)$$

Because pessimists expect greater climate damages, $E_t^P[\Delta T_t] > E_t^O[\Delta T_t]$, we have $Cap_{t+1}^P < Cap_{t+1}^O$. The median-voting cap, Cap_t^{MV} , is therefore

$$Cap_{t+1}^{MV} = \begin{cases} Cap_{t+1}^P & \text{if } \pi_t > \frac{1}{2} \\ Cap_{t+1}^O & \text{if } \pi_t < \frac{1}{2} \end{cases} \quad (53)$$

Abstracting away from the knife-edge tie case $\pi_t = \frac{1}{2}$, this can be written compactly using an indicator function:

$$Cap_{t+1}^{MV} = Cap_{t+1}^O + (Cap_{t+1}^P - Cap_{t+1}^O) \mathbf{1}_{\{\pi_t > \frac{1}{2}\}} \quad (54)$$

Having established the structure of the emissions cap under median voting, we now compare its equilibrium outcomes with those under probabilistic voting. For comparison, we denote the probabilistic-voting cap by Cap_{t+1}^{PV} , as defined in (44). Holding everything other than the voting rule fixed, the difference in portfolio structure between the two rules is driven entirely by the cap difference, $\Delta Cap_{t+1} \equiv Cap_{t+1}^{MV} - Cap_{t+1}^{PV}$. Specifically,

$$\begin{aligned} \Delta X_t^B &\equiv X_t^{B,MV} - X_t^{B,PV} = + \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} \\ \Delta X_t^G &\equiv X_t^{G,MV} - X_t^{G,PV} = - \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} \end{aligned} \quad (55)$$

In vector form, we have:

$$\Delta X_t^M \equiv X_t^{M,MV} - X_t^{M,PV} = \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} e_{BMG}. \quad (56)$$

If median voting produces a stricter cap than probabilistic voting ($\Delta Cap_{t+1} < 0$), the direct cap-clearing effect lowers the brown market-clearing weight and raises the green market-clearing weight. Conversely, if median voting produces a more lenient cap ($\Delta Cap_{t+1} > 0$), the direct cap-clearing effect raises the brown weight and lowers the green weight.

Expected return differences. We now analyze how replacing probabilistic voting with median voting affects equilibrium expected returns. The key distinction between the two mechanisms is that probabilistic voting generates a smooth policy response to the distribution of beliefs, whereas median voting generates a discontinuous response around the majority threshold.

To characterize the expected-return implications, recall that equilibrium expected excess returns are given by (37). A full comparison of the two rules must account for changes in both market allocations and portfolio incentives. First, the voting rule changes the cap-clearing market portfolio X_t^M . Second, because the policy change alters the portfolio problem of each type, it may also change the type-specific portfolios X_t^O and X_t^P .

For a fixed belief distribution π_t and fixed subjective climate beliefs $\mu_{\Delta T}^P$ and $\mu_{\Delta T}^O$, we define the difference

in the BMG expected return across the two voting rules as

$$\Delta\mu_{r,t}^{BMG} = \mu_{r,t}^{BMG}(Cap_{t+1}^{MV}) - \mu_{r,t}^{BMG}(Cap_{t+1}^{PV}), \quad (57)$$

A closed-form expression for (57) is not available in general. Under one simplifying assumption, however, we can obtain a closed form and interpret it. Recall the valuation channel presented earlier, and suppose that an innovation in cumulative emissions contributes equally to brown and green returns through future valuations. Formally,

$$\rho_B B_{B,E} = \rho_G B_{G,E} \equiv \nu_t \quad (58)$$

The purpose of this assumption is to shut down the differential valuation channel through which cumulative emissions affect brown returns more strongly than green returns. Under (58), the difference in the BMG expected return under median voting relative to probabilistic voting is

$$\Delta\mu_{r,t}^{BMG} = \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} \{ \mathcal{S}_t + \nu_t [\gamma(1 - \alpha) + (1 - \gamma)b \{ \hat{\pi}_t \pi_t + (1 - \hat{\pi}_t)(1 - \pi_t) \}] e'_{BMG} \Sigma_{\Delta T} e_{BMG} \} \quad (59)$$

$$\text{where } \hat{\pi}_t = \frac{\pi_t \mathcal{A}_t^O - (1 - \pi_t) \mathcal{J}_t^O}{e'_{BMG} \mathcal{D} e_{BMG}}$$

By (59), the BMG expected-return differential across the two rules depends on which type—pessimist or optimist—is in the majority and on the sign of the term in braces. Table 2 summarizes the result.

Table 2. Median vs. probabilistic voting caps and the BMG difference

Expression in braces in (59)	Stricter median-voting cap	More lenient median-voting cap
Positive	$\Delta\mu_{r,t}^{BMG} < 0$	$\Delta\mu_{r,t}^{BMG} > 0$
Negative	$\Delta\mu_{r,t}^{BMG} > 0$	$\Delta\mu_{r,t}^{BMG} < 0$
Zero	No BMG difference	No BMG difference

To understand the intuition, it is useful to separate the term in braces into two parts. The first, $\mathcal{S}_t$, captures the sensitivity of the BMG compensation to the brown market weight, holding portfolio coefficients fixed. The second,

$$\nu_t [\gamma(1 - \alpha) + (1 - \gamma)b \{ \hat{\pi}_t \pi_t + (1 - \hat{\pi}_t)(1 - \pi_t) \}] e'_{BMG} \Sigma_{\Delta T} e_{BMG}$$

is the valuation-related adjustment to that sensitivity, which accounts for climate exposure, risk aversion, and damage internalization.

The result therefore says that, under a more lenient (resp. stricter) median-voting cap, the difference in the BMG expected return is positive (resp. negative) if and only if the additional compensation required to absorb more brown exposure exceeds the reduction in compensation produced by the valuation adjustment.

5 Conclusion

This paper examines how disagreement about the economic damages from climate change shapes asset prices and the performance of sustainable investing strategies. Standard climate asset-pricing models predict that brown firms should offer higher expected returns than green firms as compensation for their greater exposure to climate and transition risks, implying that a brown-minus-green (BMG) portfolio should earn a positive expected return. However, empirical studies document mixed evidence on the expected BMG returns: some argue that they follow a time-varying pattern, while others show that the sign can change over time. This disconnect between theoretical predictions and empirical evidence motivates our study.

We propose that belief dispersion—heterogeneity in investors’ expectations about climate damages and the transition path—is an important state variable that existing models have largely overlooked and that can help explain this disconnect. We first provide suggestive evidence using institutional portfolios as a market-based proxy for climate beliefs. Specifically, we measure belief dispersion using variation in asset managers’ exposure to firms’ carbon emissions that cannot be explained by standard portfolio characteristics, manager type, or investment strategy. This measure varies substantially over time, and periods of greater dispersion are followed by lower BMG returns. Although suggestive, this evidence points to a channel through which disagreement about climate damages and future climate policy can affect asset prices through distorted risk-sharing and speculative trading.

Motivated by these patterns, we develop an overlapping-generations equilibrium model in which purely pecuniary investors internalize climate damages but disagree about their magnitude. Asset prices are determined in equilibrium, and valuation ratios respond to cumulative emissions and climate beliefs. On top of the standard two-factor structure with financial and climate risk premia, internalizing damages under subjective beliefs gives rise to speculation on climate beliefs, distorted risk sharing, and a belief-driven valuation wedge. These additional pricing forces can make the expected BMG premium positive, zero, or negative, depending on the degree of belief dispersion, the political pressure that determines the stringency of the emissions cap, and the

valuation state. Under certain conditions derived in the model, sufficiently convex climate damages restore a positive BMG premium. Allowing beliefs to evolve endogenously also generates a conditional prediction of mean reversion in the BMG premium.

The empirical analysis also has a number of limitations. Portfolio tilts may capture mandates, benchmark choices, investment constraints, or risk-management practices in addition to investors' beliefs. The relatively short quarterly sample also limits statistical power, and the distinction between brown and green firms depends on how emissions are measured and on data coverage. Future work could strengthen the holdings-based measure by comparing it with proxy voting and manager disclosures, examine the mechanism in countries with different transition-policy regimes, and study whether financing constraints or limits to arbitrage make belief-driven price effects stronger.

Despite these limitations, the results point to an important implication. As long as climate beliefs remain polarized, sustainable-investing premia are likely to depend on the distribution of beliefs, not just with underlying physical and regulatory risks. Disagreement is therefore a potential candidate for explaining why brown vs. green return spreads are unstable and why estimates of climate risk premia can differ across periods and settings.

References

- Acharya, V. V. and L. H. Pedersen (2005). Asset pricing with liquidity risk. *Journal of financial Economics* 77(2), 375–410.
- Albuquerque, R., Y. Koskinen, and C. Zhang (2019). Corporate social responsibility and firm risk: Theory and empirical evidence. *Management science* 65(10), 4451–4469.
- Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2023). Climate change concerns and the performance of green vs. brown stocks. *Management Science* 69(12), 7607–7632.
- Atilgan, Y., K. O. Demirtas, A. Edmans, and A. D. Gunaydin (2024). Does the carbon premium reflect risk or mispricing? *FEB-RN Research Paper* (03).
- Atmaz, A. and S. Basak (2018). Belief dispersion in the stock market. *The Journal of Finance* 73(3), 1225–1279.
- Avramov, D., A. Lioui, Y. Liu, and A. Tarelli (2025). Dynamic esg equilibrium. *Management Science* 71(4), 2867–2889.
- Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics* 131(4), 1593–1636.
- Baldauf, M., L. Garlappi, and C. Yannelis (2020). Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies* 33(3), 1256–1295.
- Barnett, M. (2024). A run on fossil fuel? climate change and transition risk. *arXiv preprint arXiv:2410.00902*.
- Barnett, M., W. Brock, and L. P. Hansen (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies* 33(3), 1024–1066.
- Basak, S. (2005). Asset pricing with heterogeneous beliefs. *Journal of Banking & Finance* 29(11), 2849–2881.
- Bolton, P. and M. Kacperczyk (2021). Do investors care about carbon risk? *Journal of financial economics* 142(2), 517–549.
- Bolton, P. and M. Kacperczyk (2023). Global pricing of carbon-transition risk. *The Journal of Finance* 78(6), 3677–3754.

- Brock, W. A. and C. H. Hommes (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control* 22(8-9), 1235–1274.
- Campbell, J. Y. and R. J. Shiller (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The review of financial studies* 1(3), 195–228.
- Campbell, J. Y. and L. M. Viceira (2002). *Strategic asset allocation: portfolio choice for long-term investors*. OUP Oxford.
- Ceccarelli, M. and S. Ramelli (2024). Climate transition beliefs. *Swiss Finance Institute Research Paper* (24-22).
- David, A. (2008). Heterogeneous beliefs, speculation, and the equity premium. *The Journal of Finance* 63(1), 41–83.
- De Angelis, T., P. Tankov, and O. D. Zerbib (2023). Climate impact investing. *Management Science* 69(12), 7669–7692.
- Dumas, B., A. Kurshev, and R. Uppal (2009). Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. *The Journal of Finance* 64(2), 579–629.
- Garleanu, N. and L. H. Pedersen (2025). Climate risk pricing. *Available at SSRN* 5978854.
- Gavriilidis, K., D. R. Känzig, R. Raghavan, and J. H. Stock (2026). The macroeconomic effects of climate policy uncertainty. Technical Report 34762, National Bureau of Economic Research.
- Giglio, S., M. Maggiori, J. Stroebe, Z. Tan, S. Utkus, and X. Xu (2025). Four facts about esg beliefs and investor portfolios. *Journal of Financial Economics* 164, 103984.
- Hambel, C., H. Kraft, and E. Schwartz (2021). Optimal carbon abatement in a stochastic equilibrium model with climate change. *European Economic Review* 132, 103642.
- Harrison, J. M. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics* 92(2), 323–336.
- Heinkel, R., A. Kraus, and J. Zechner (2001). The effect of green investment on corporate behavior. *Journal of Financial and Quantitative Analysis* 36(4), 431–449.

- Hong, H. and D. A. Sraer (2016). Speculative betas. *The Journal of Finance* 71(5), 2095–2144.
- Hong, H., N. Wang, and J. Yang (2023). Mitigating disaster risks in the age of climate change. *Econometrica* 91(5), 1763–1802.
- Hsu, P.-H., K. Li, and C.-Y. Tsou (2023). The pollution premium. *The Journal of Finance* 78(3), 1343–1392.
- Ilhan, E., Z. Sautner, and G. Vilkov (2021). Carbon tail risk. *The Review of Financial Studies* 34(3), 1540–1571.
- Jagannathan, R., I. Meier, and V. Sokolovski (2025). Dirty business: Transition risk of factor portfolios. Technical report, National Bureau of Economic Research.
- Jouini, E. and C. Napp (2006). Heterogeneous beliefs and asset pricing in discrete time: An analysis of pessimism and doubt. *Journal of Economic Dynamics and Control* 30(7), 1233–1260.
- Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009). The proportionality of global warming to cumulative carbon emissions. *Nature* 459(7248), 829–832.
- Pástor, L., R. F. Stambaugh, and L. A. Taylor (2021). Sustainable investing in equilibrium. *Journal of financial economics* 142(2), 550–571.
- Pástor, L., R. F. Stambaugh, and L. A. Taylor (2022). Dissecting green returns. *Journal of Financial Economics* 146(2), 403–424.
- Pedersen, L. H., S. Fitzgibbons, and L. Pomorski (2021). Responsible investing: The esg-efficient frontier. *Journal of financial economics* 142(2), 572–597.
- Sautner, Z., L. Van Lent, G. Vilkov, and R. Zhang (2023). Pricing climate change exposure. *Management Science* 69(12), 7540–7561.
- Weitzman, M. L. (2012). Ghg targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory* 14(2), 221–244.
- Zerbib, O. D. (2022). A sustainable capital asset pricing model (s-capm): Evidence from environmental integration and sin stock exclusion. *Review of Finance* 26(6), 1345–1388.
- Zhang, S. (2025). Carbon returns across the globe. *The Journal of Finance* 80(1), 615–645.

Appendix

Empirical Appendix

This appendix reports the empirical specifications omitted from the main text. Table 3 repeats the main text analysis using a 50% portfolio emissions-coverage threshold.

Table 3. Residual dispersion at the 50% coverage threshold

	(1)	(2)	(3)
Residual dispersion	-0.844* (-1.85)	-0.725* (-1.87)	-1.048*** (-2.90)
Observations	39	39	39
R^2	0.059	0.437	0.506
Circular-shift p -value	0.103	0.103	0.128
FF5 controls	No	Yes	Yes
VIX and PCI controls	No	No	Yes

Notes: This table repeats the main specifications after retaining manager-quarters with at least 50% emissions coverage.

Disagreement and mean tilt are standardized within the sample. Newey–West t -statistics with four lags are in parentheses.

All columns within a panel use the same sample.

Validation of the Disagreement Proxy

This appendix reports the two validation exercises summarized in Section 2.2.1. Both are estimated on the 2014Q1–2023Q3 sample of 39 quarters, and all variables are standardized to zero mean and unit variance within that sample. So, coefficients should be read as standard-deviation responses. We use Newey–West standard errors with four quarterly lags.

Convergent Validity: Climate-Policy Uncertainty

We ask whether residual institutional dispersion moves with an independent measure of the same underlying phenomenon. Our measure is the news-based climate-policy uncertainty index of Gavrilidis et al. (2026), constructed from newspaper text and therefore sharing neither input data nor methodology with our holdings-

based proxy. We average the monthly index values within quarter and estimate

$$U_t^{\text{climate},k} = \alpha_k + \beta_k D_t + \varepsilon_t^k, \quad (\text{A.1})$$

where $U_t^{\text{climate},k}$ is the standardized log value of index variant k in quarter t and D_t is standardized residual dispersion as defined in (1). The coefficient of interest is β_k : a positive value indicates that institutional carbon positions are more dispersed in quarters when climate policy is more contested in public debate, which is what the proxy is intended to capture. We estimate (A.1) separately for the narrow, broad-dictionary, and legacy variants of the index rather than selecting one, since the three differ in the breadth of their keyword dictionaries and no variant is obviously preferred for our purpose.

Table 4 reports the results. The estimated β_k ranges from 0.574 to 0.648 and is significant at the 1% level in all three specifications, with R^2 between 0.33 and 0.42.

Table 4. Residual institutional dispersion and climate-policy uncertainty

	Narrow	Broad	Legacy
	(1)	(2)	(3)
Residual dispersion	0.622*** (5.22)	0.648*** (6.19)	0.574*** (4.84)
Constant	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)
Observations	39	39	39
R^2	0.387	0.420	0.330

Notes: The dependent variables are the standardized log narrow, broad, and legacy news-based climate-policy uncertainty indexes, respectively. Monthly index values are averaged within quarter. Residual dispersion is standardized within the 2014Q1–2023Q3 estimation sample. Newey–West t -statistics with four quarterly lags are in parentheses. *** denotes significance at the 1% level.

The relation is not robust to detrending. Both series rise over a decade in which climate policy became steadily more salient, and after including a linear trend the coefficient in the narrow-index specification falls to -0.085 ($t = -0.73$). With 39 quarters we cannot separate a common trend from higher-frequency comovement, and we therefore interpret the result as evidence that the proxy loads on climate-policy contestation at medium frequencies rather than as a quarter-to-quarter relationship.

Discriminant Validity: Placebo Test

A natural competing interpretation is that dispersed carbon positioning reflects generalized uncertainty rather than anything specific to climate: when managers disagree about the macroeconomic outlook, they disagree about carbon exposure as a by-product. Under that reading D_t is a repackaged macro-uncertainty proxy, and any relation to future returns would carry no climate content.

We test this by replacing the dependent variable with a measure of aggregate uncertainty that is not climate-specific, the broad U.S. Economic Policy Uncertainty Index of Baker et al. (2016), and conditioning on climate-policy uncertainty:

$$U_t^{\text{general}} = \alpha + \beta D_t + \delta U_t^{\text{climate}} + \varepsilon_t, \quad (\text{A.2})$$

where U_t^{general} is the standardized log EPU index averaged within quarter. The test has a clear prediction in each direction. If D_t captures economy-wide uncertainty, it should retain explanatory power for the placebo outcome even after the climate component is partialled out, so β should remain positive and significant. If instead D_t is specific to climate, its apparent relation to general uncertainty should be absorbed by U_t^{climate} , leaving β indistinguishable from zero.

Table 5 reports the estimates. Column 1 omits the climate control: the unconditional relation between dispersion and general policy uncertainty is 0.260 and significant only at the 10% level, already weaker than the 0.57–0.65 range in Table 4. Columns 2 and 3 add the narrow and broad climate-policy uncertainty indexes respectively. In both, the coefficient on residual dispersion turns negative and statistically insignificant, while climate-policy uncertainty remains strongly related to the placebo outcome. The dispersion proxy therefore loads on the climate-policy component of aggregate uncertainty rather than on aggregate uncertainty as such.

Table 5. Placebo test using general economic-policy uncertainty

	(1)	(2)	(3)
Residual dispersion	0.260*	-0.156	-0.264
	(1.91)	(-0.90)	(-1.50)
Climate-policy uncertainty		0.669***	0.808***
		(3.53)	(4.06)
Constant	0.000	0.000	0.000
	(0.00)	(0.00)	(0.00)
Climate index	None	Narrow	Broad
Observations	39	39	39
R^2	0.067	0.342	0.446

Notes: The dependent variable is the standardized log U.S. Economic Policy Uncertainty Index (FRED series USEPUINDXM), averaged within quarter. Columns 2 and 3 control for the standardized log narrow and broad climate-policy uncertainty indexes, respectively. Residual dispersion is standardized within the 2014Q1–2023Q3 estimation sample. Newey–West t -statistics with four quarterly lags are in parentheses. *** and * denote significance at the 1% and 10% levels.

Two caveats apply. The 39-quarter sample gives the falsification exercise limited power, so failure to reject is weaker evidence than it would be in a longer sample. And climate-policy uncertainty is itself a component of general policy uncertainty, so the two right-hand-side variables in (A.2) are correlated by construction; the test establishes that dispersion carries no incremental information about the non-climate component, not that the two components are cleanly separable.

B.1 Derivation of the Reduced Portfolio Problem and Equilibrium Expected Returns

Suppose that investors enter with common wealth. So the aggregation weights are their population shares.

Write $\pi_t^P = \pi_t$, $\pi_t^O = 1 - \pi_t$ and

$$X_t = \pi_t X_t^P + (1 - \pi_t) X_t^O, \quad X_t^i = e_G + X_t^{i,B} e_{BMG}, \quad e_G = (0, 1)', \quad e_{BMG} = (1, -1)'.$$

Both types evaluate aggregate temperature. A group's reallocation changes the aggregate portfolio by $\pi_t^i e_{BMG}$ per unit change in its brown weight.

Define the vector of expected log temperature contributions as

$$\mu_{\Delta T}^i \equiv E_t^i[\log \Delta T_t]. \quad (\text{B.1})$$

Using the first-order approximation $\log \Delta T_t \simeq X_t' \log \Delta T_t$, we obtain

$$E_t^i[\log \Delta T_t] \simeq X_t' \mu_{\Delta T}^i, \quad \log \Delta T_t - E_t^i[\log \Delta T_t] \simeq X_t' \{u_{t+1} - E_t^i[u_{t+1}]\}, \quad (\text{B.2})$$

where u_{t+1} is the vector of log temperature-contribution innovations under the objective measure. The approximation is taken around a common reference value of the two asset-level temperature contributions.

Quadratic objective. Let y_t^i denote log damage-adjusted terminal wealth. Under the Gaussian approximation, maximizing expected CRRA utility is equivalent to maximizing

$$E_t^i[y_t^i] + \frac{1-\gamma}{2} \text{Var}_t^i(y_t^i),$$

with the log-utility limit at $\gamma = 1$. Apply the log-portfolio approximation and the damage linearization with elasticity

$$b = \frac{\eta \delta \overline{\Delta T}^\eta}{1 + \delta \overline{\Delta T}^\eta}.$$

The own-return variance terms combine to $-\gamma X_t^{i'} \Sigma_{r,t} X_t^i / 2$. The covariance between log returns and log damages contributes $-(1-\gamma)b X_t^{i'} \Sigma_{rT,t} X_t^i$, and damage variance contributes $(1-\gamma)b^2 X_t^{i'} \Sigma_{\Delta T} X_t^i / 2$. Therefore, the terms relevant to portfolio choice are

$$\begin{aligned} & X_t^{i'} (E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{r,t}^2) - \frac{\gamma}{2} X_t^{i'} \Sigma_{r,t} X_t^i - b X_t^{i'} \mu_{\Delta T}^i \\ & - (1-\gamma)b X_t^{i'} \Sigma_{rT,t} X_t^i + \frac{1-\gamma}{2} b^2 X_t^{i'} \Sigma_{\Delta T} X_t^i. \end{aligned} \quad (\text{B.3})$$

Substituting the population-weighted aggregate and collecting the own-portfolio and interaction terms gives, up to terms independent of X_t^i ,

$$X_t^{i'} a_t^i - \frac{\gamma}{2} X_t^{i'} \Sigma_{\text{Total},t}^i X_t^i - X_t^{i'} \mathcal{K}_t^i X_t^{i-}, \quad (\text{B.4})$$

where

$$\begin{aligned}
a_t^i &= E_t^i[\mathbf{r}_{t+1}] - r_f \mathbf{1} + \frac{1}{2} \sigma_{r,t}^2 - b \pi_t^i \mu_{\Delta T}^i, \\
\gamma \Sigma_{\text{Total},t}^i &= \gamma \Sigma_{r,t} + (1 - \gamma) b \pi_t^i (\Sigma_{rT,t} + \Sigma'_{rT,t}) - (1 - \gamma) b^2 (\pi_t^i)^2 \Sigma_{\Delta T}, \\
\mathcal{K}_t^i &= (1 - \gamma) b \pi_t^{i-} (\Sigma_{rT,t} - b \pi_t^i \Sigma_{\Delta T}).
\end{aligned} \tag{B.5}$$

The return covariance is common across types. The matrix $\gamma \Sigma_{\text{Total},t}^i$ measures the curvature of each type's damage-adjusted objective.

Return moments and beliefs. Maintain the local innovation specification

$$\Delta d_{t+1} - E_t[\Delta d_{t+1}] = \alpha \varepsilon_{K,t+1} + (1 - \alpha) u_{t+1}, \quad Emission_{t+1}^{cum} - E_t[Emission_{t+1}^{cum}] = X_t^{M'} u_{t+1}.$$

The capital innovation is orthogonal to u_{t+1} . The emissions and temperature innovations are expressed in the common local units implicit in these loadings. Under the timing and two-state specification stated in B.7–B.8, π_{t+1} is known at date t . Hence

$$pd_{t+1}^j - E_t[pd_{t+1}^j] = B_{j,E} X_t^{M'} u_{t+1}, \tag{B.6}$$

$$\begin{aligned}
L_{r,t} &= (1 - \alpha) I + \rho B_E X_t^{M'}, \\
\Sigma_{r,t} &= \alpha^2 \Sigma_K + L_{r,t} \Sigma_{\Delta T} L'_{r,t}, \quad \Sigma_{rT,t} = L_{r,t} \Sigma_{\Delta T}.
\end{aligned} \tag{B.7}$$

Valuation news incorporates revisions to expected future dividends and required compensation.

Let the scalar tilted innovation $Z_{\Delta T}$ be standard Gaussian, jointly Gaussian with u_{t+1} , and let $\sigma_{\Delta T} = \text{Cov}_t(u_{t+1}, Z_{\Delta T})$. Assume $\Sigma_{\Delta T} - \sigma_{\Delta T} \sigma'_{\Delta T}$ is positive semidefinite. Under the density $\exp(\theta_i Z_{\Delta T} - \theta_i^2/2)$, the Gaussian mean shifts and its covariance is unchanged. It follows that

$$\mu_{\Delta T}^i = \mu_{\Delta T} + \theta_i \sigma_{\Delta T}, \quad E_t^i[\mathbf{r}_{t+1}] = E_t[\mathbf{r}_{t+1}] + \theta_i L_{r,t} \sigma_{\Delta T}. \tag{B.8}$$

The common-factor benchmark used in B.5 specializes this specification to $u_{t+1} = \sigma_{\Delta T} Z_{\Delta T}$.

Proof of Proposition 1. For the vector solution and asset-level pricing, maintain the following full-investment compatibility condition: at the equilibrium return vector and common r_f , both unconstrained maximizers of

(B.4) are strictly positive and sum to one. This selects the class of interior equilibria to which the displayed vector solution applies. Assume the curvature matrices are positive definite and $\rho(M_t^O M_t^P) < 1$.

Strict concavity makes the vector first-order conditions sufficient:

$$\gamma \Sigma_{\text{Total},t}^i X_t^i + \mathcal{K}_t^i X_t^{i-} = a_t^i. \quad (\text{B.9})$$

Premultiplication by $(\gamma \Sigma_{\text{Total},t}^i)^{-1}$ yields the best responses in the main text. Substitute one into the other. Using the paper's $M_t^i = (\gamma \Sigma_{\text{Total},t}^i)^{-1} \mathcal{K}_t^i$, this gives

$$\begin{aligned} X_t^P &= (I - M_t^P M_t^O)^{-1} [(\gamma \Sigma_{\text{Total},t}^P)^{-1} a_t^P - M_t^P (\gamma \Sigma_{\text{Total},t}^O)^{-1} a_t^O], \\ X_t^O &= (I - M_t^O M_t^P)^{-1} [(\gamma \Sigma_{\text{Total},t}^O)^{-1} a_t^O - M_t^O (\gamma \Sigma_{\text{Total},t}^P)^{-1} a_t^P]. \end{aligned} \quad (\text{B.10})$$

The spectral-radius restriction makes both inverses well defined by their Neumann series. It also makes successive two-step best responses converge. The compatibility condition places these maximizers on the fully invested feasible set, proving the stated interior portfolio solution.

For allocation calculations, impose full investment directly by substitution. The resulting first-order condition is

$$e'_{BMG} [a_t^i - \gamma \Sigma_{\text{Total},t}^i X_t^i - \mathcal{K}_t^i X_t^{i-}] = 0. \quad (\text{B.11})$$

With $\mathcal{A}_t^i = \gamma e'_{BMG} \Sigma_{\text{Total},t}^i e_{BMG}$ and $\mathcal{J}_t^i = e'_{BMG} \mathcal{K}_t^i e_{BMG}$, this gives

$$X_t^{i,B} = \frac{e'_{BMG} a_t^i - e'_{BMG} (\gamma \Sigma_{\text{Total},t}^i + \mathcal{K}_t^i) e_G}{\mathcal{A}_t^i} - \frac{\mathcal{J}_t^i}{\mathcal{A}_t^i} X_t^{i-,B}. \quad (\text{B.12})$$

Its response is offsetting when $\mathcal{J}_t^i > 0$ and reinforcing when $\mathcal{J}_t^i < 0$, provided $\mathcal{A}_t^i > 0$. At log utility the interaction is zero. These scalar allocation equations will also be used for comparisons that condition on full investment directly.

Demand gap. Market clearing implies

$$X_t^P = X_t^M - (1 - \pi_t) \Delta X_t e_{BMG}, \quad X_t^O = X_t^M + \pi_t \Delta X_t e_{BMG}. \quad (\text{B.13})$$

Subtract the two conditions (B.11), use (B.8), and substitute (B.13). The result is

$$\begin{aligned}
\mathcal{D} &= \gamma(1 - \pi_t)\Sigma_{\text{Total},t}^P + \gamma\pi_t\Sigma_{\text{Total},t}^O - \pi_t\mathcal{K}_t^P - (1 - \pi_t)\mathcal{K}_t^O, \\
\mathcal{N} &= b[\pi_t\mu_{\Delta T}^P - (1 - \pi_t)\mu_{\Delta T}^O] - (\theta_P - \theta_O)L_{r,t}\sigma_{\Delta T} \\
&\quad + [\gamma\Sigma_{\text{Total},t}^P + \mathcal{K}_t^P - \gamma\Sigma_{\text{Total},t}^O - \mathcal{K}_t^O]X_t^M, \\
X_t^O - X_t^P &= \frac{e'_{BMG}\mathcal{N}}{e'_{BMG}\mathcal{D}e_{BMG}}e_{BMG} = \Delta X_t e_{BMG}.
\end{aligned} \tag{B.14}$$

Assume the denominator is nonzero and the resulting weights are interior. Whenever the holding order in Lemma 1 is used, maintain $e'_{BMG}\mathcal{N}/(e'_{BMG}\mathcal{D}e_{BMG}) > 0$.

Equilibrium expected returns. Insert (B.8) into (B.9), population-weight the two equations, and use market clearing. Under the full-investment compatibility condition,

$$\mu_{r,t} = \gamma\Sigma_{r,t}X_t^M + b[\pi_t^2\mu_{\Delta T}^P + (1 - \pi_t)^2\mu_{\Delta T}^O] + Hedging_t - \tilde{\theta}_t L_{r,t}\sigma_{\Delta T}, \tag{B.15}$$

$$\begin{aligned}
Hedging_t &= (1 - \gamma)b\left\{\Sigma_{rT,t}X_t^M + \Sigma'_{rT,t}[\pi_t^2X_t^P + (1 - \pi_t)^2X_t^O] \right. \\
&\quad \left. - b[\pi_t^2 + (1 - \pi_t)^2]\Sigma_{\Delta T}X_t^M\right\}.
\end{aligned} \tag{B.16}$$

To see the simplification, the coefficient of $\Sigma_{rT,t}$ is $\pi_t^2X_t^P + (1 - \pi_t)^2X_t^O + \pi_t(1 - \pi_t)(X_t^P + X_t^O) = X_t^M$. The coefficient of the transpose is $\pi_t^2X_t^P + (1 - \pi_t)^2X_t^O$, and that of $-b\Sigma_{\Delta T}$ is $[\pi_t^2 + (1 - \pi_t)^2]X_t^M$. This proves the expected-return decomposition. Premultiplication by e'_{BMG} gives the BMG equation directly from the allocation conditions.

For conditionally lognormal returns, $\mu_{r,t}^j = \log E_t[R_{t+1}^j] - \log R_f$. Thus $\mu_{r,t}^{BMG}$ is the log ratio of the two expected gross returns and has the same sign as their arithmetic difference. The comparative statics use the paper's variance-adjusted log-return measure. At log utility $Hedging_t = 0$. Under consensus the mean-damage and hedging terms take their common-belief values. ■

B.2 Proof of Proposition 2 (Factor Representation)

Maintain the return-moment specification and full-investment compatibility condition in B.1. Define the financial and climate factors as $\alpha X_t^{M'}\varepsilon_{K,t+1}$ and $X_t^{M'}L_{r,t}u_{t+1}$. Assume both have positive conditional variance.

Orthogonality of the primitive innovations makes the factors orthogonal. Their covariance betas and risk prices are

$$\begin{aligned}\beta_{F,t}^j &= \frac{[\alpha^2 \Sigma_K X_t^M]_j}{X_t^{M'} \alpha^2 \Sigma_K X_t^M}, & \lambda_{F,t} &= \gamma X_t^{M'} \alpha^2 \Sigma_K X_t^M, \\ \beta_{C,t}^j &= \frac{[L_{r,t} \Sigma_{\Delta T} L'_{r,t} X_t^M]_j}{X_t^{M'} L_{r,t} \Sigma_{\Delta T} L'_{r,t} X_t^M}, & \lambda_{C,t} &= \gamma X_t^{M'} L_{r,t} \Sigma_{\Delta T} L'_{r,t} X_t^M.\end{aligned}\tag{B.17}$$

Multiplying each beta by its risk price recovers the corresponding component of $\gamma \Sigma_{r,t} X_t^M$. Define

$$Speculation_t^j = [\pi_t^2 \mu_{\Delta T}^P + (1 - \pi_t)^2 \mu_{\Delta T}^O]_j, \quad Valuation_t^j = -\tilde{\theta}_t [L_{r,t} \sigma_{\Delta T}]_j.$$

Taking component j of (B.15) therefore gives

$$\mu_{r,t}^j = \beta_{F,t}^j \lambda_{F,t} + \beta_{C,t}^j \lambda_{C,t} + b Speculation_t^j + Hedging_t^j + Valuation_t^j.\tag{B.18}$$

This proves the proposition. The climate-factor price is nonnegative, and it is strictly positive under the positive-variance assumption. In the proportional-loading benchmark $L_{r,t} \Sigma_{\Delta T} = \varsigma \Sigma_{\Delta T}$,

$$\lambda_{C,t} = \gamma \varsigma^2 X_t^{M'} \Sigma_{\Delta T} X_t^M.$$

The factor prices ordinary climate-related return risk; the damage adjustments are recorded separately in the remaining pricing terms. The standard two-factor benchmark obtains when these additional adjustments are zero, as under $b = 0$ and zero return-belief tilt. ■

B.3 Proof of Proposition 4

Consider a differentiable interior equilibrium. Hold b , each type's temperature mean and probability tilt, the reference emissions exposures, and primitive covariance matrices fixed in this local comparison. The level expectations entering political pressure are fixed inputs to the comparison. Return loadings and portfolio coefficients respond as specified in G_t^i . Assume $Emission_{t+1}^B > Emission_{t+1}^G$, $Emission_{t+1}^G < Cap_{t+1} < Emission_{t+1}^B$, and $\mathcal{S}_t > 0$ with nonzero denominators.

The binding cap and full investment imply

$$X_t^{M,B} = \frac{Cap_{t+1} - Emission_{t+1}^G}{Emission_{t+1}^B - Emission_{t+1}^G}, \quad \frac{dX_t^{M,B}}{d\pi_t} = -\frac{\psi Cap_{t+1}}{Emission_{t+1}^B - Emission_{t+1}^G} \frac{d\phi_t}{d\pi_t}. \quad (\text{B.19})$$

The pressure derivative, holding each type's level expectation fixed, is

$$\frac{d\phi_t}{d\pi_t} = \frac{E_t^P[\Delta T_t] + E_t^O[\Delta T_t]}{\Delta T}.$$

Positive expected warming makes this derivative positive. Differentiating each projected first-order condition gives

$$\begin{aligned} \frac{d\mu_{r,t}^{BMG}}{d\pi_t} &= \mathcal{A}_t^P \frac{dX_t^{P,B}}{d\pi_t} + \mathcal{J}_t^P \frac{dX_t^{O,B}}{d\pi_t} + G_t^P, \\ \frac{d\mu_{r,t}^{BMG}}{d\pi_t} &= \mathcal{J}_t^O \frac{dX_t^{P,B}}{d\pi_t} + \mathcal{A}_t^O \frac{dX_t^{O,B}}{d\pi_t} + G_t^O. \end{aligned} \quad (\text{B.20})$$

Here the coefficient-adjustment term is

$$G_t^i = e'_{BMG} \left[\gamma \frac{d\Sigma_{\text{Total},t}^i}{d\pi_t} X_t^i + \frac{d\mathcal{K}_t^i}{d\pi_t} X_t^{i-} + b \frac{d\pi_t^i}{d\pi_t} \mu_{\Delta T}^i - \theta_i \frac{dL_{r,t}}{d\pi_t} \sigma_{\Delta T} \right].$$

Differentiating market clearing gives

$$\pi_t \frac{dX_t^{P,B}}{d\pi_t} + (1 - \pi_t) \frac{dX_t^{O,B}}{d\pi_t} = \frac{dX_t^{M,B}}{d\pi_t} + \Delta X_t. \quad (\text{B.21})$$

Inverting the two-by-two coefficient matrix in (B.20) and substituting into (B.21) yields

$$\frac{d\mu_{r,t}^{BMG}}{d\pi_t} = \mathcal{S}_t \left[\Delta X_t - \frac{\psi Cap_{t+1}}{Emission_{t+1}^B - Emission_{t+1}^G} \frac{d\phi_t}{d\pi_t} \right] + \mathcal{R}_t, \quad (\text{B.22})$$

where direct inversion gives

$$\begin{aligned} \mathcal{S}_t &= \frac{\mathcal{A}_t^P \mathcal{A}_t^O - \mathcal{J}_t^P \mathcal{J}_t^O}{e'_{BMG} \mathcal{D}e_{BMG}}, \\ \mathcal{R}_t &= \frac{[\pi_t \mathcal{A}_t^O - (1 - \pi_t) \mathcal{J}_t^O] G_t^P + [(1 - \pi_t) \mathcal{A}_t^P - \pi_t \mathcal{J}_t^P] G_t^O}{e'_{BMG} \mathcal{D}e_{BMG}}. \end{aligned}$$

In particular,

$$e'_{BMG} \mathcal{D}e_{BMG} = (1 - \pi_t) \mathcal{A}_t^P + \pi_t \mathcal{A}_t^O - \pi_t \mathcal{J}_t^P - (1 - \pi_t) \mathcal{J}_t^O.$$

The assumed positive slope converts the compensation comparison into a comparison of brown portfolio weights. Divide (B.22) by $d\phi_t/d\pi_t > 0$. The BMG premium rises with political pressure along this mechanism if and only if

$$\Delta X_t + \frac{\mathcal{R}_t}{\mathcal{S}_t} > \frac{\psi Cap_{t+1}}{Emission_{t+1}^B - Emission_{t+1}^G} \frac{d\phi_t}{d\pi_t}. \quad (\text{B.23})$$

The left side expresses the demand response, including changes in portfolio incentives, in units of brown portfolio weight. The right side is the decline in the brown market weight implied by the cap. Equality gives no first-order BMG response. The reverse inequality gives a decline. For the separate asset-level signs, additionally hold the common compensation component fixed along the pressure comparison, $d(\mu_{r,t}^B + \mu_{r,t}^G)/d\phi_t = 0$.

Then

$$\frac{d\mu_{r,t}^B}{d\phi_t} = \frac{1}{2} \frac{d\mu_{r,t}^{BMG}}{d\phi_t}, \quad \frac{d\mu_{r,t}^G}{d\phi_t} = -\frac{1}{2} \frac{d\mu_{r,t}^{BMG}}{d\phi_t}.$$

Under the strict inequality in (B.23), brown compensation rises and green compensation falls. All coefficient derivatives must be taken along the stated comparison. For example, at fixed ρB_E ,

$$\begin{aligned} \frac{dL_{r,t}}{d\pi_t} &= \rho B_E e'_{BMG} \frac{dX_t^{M,B}}{d\pi_t}, \quad \frac{d\Sigma_{rT,t}}{d\pi_t} = \frac{dL_{r,t}}{d\pi_t} \Sigma_{\Delta T}, \\ \frac{d\Sigma_{r,t}}{d\pi_t} &= \frac{dL_{r,t}}{d\pi_t} \Sigma_{\Delta T} L'_{r,t} + L_{r,t} \Sigma_{\Delta T} \frac{dL'_{r,t}}{d\pi_t}. \end{aligned}$$

Differentiating (B.5) then supplies the derivatives in G_t^i . The derivatives in G_t^i are evaluated under these maintained local restrictions. ■

B.4 Proof of Property 1

Parameter region. Assume $\theta_O \leq 0$, $\mathcal{B}_t^O > 0$, and feasible interior portfolios throughout the comparison. The affine belief decomposition is a fixed-coefficient comparison: hold π_t , X_t^M , b , $L_{r,t}$ and the conditional covariance matrices fixed while varying the two probability tilts. The following calculation also gives an explicit expression for BMG after eliminating the individual portfolios.

Define the effective weight

$$\hat{\pi}_t = \frac{\pi_t \mathcal{A}_t^O - (1 - \pi_t) \mathcal{J}_t^O}{e'_{BMG} \mathcal{D} e_{BMG}}. \quad (\text{B.24})$$

This algebraic weight eliminates the holding gap from the two pricing equations. It satisfies

$$\hat{\pi}_t [(1 - \pi_t) \mathcal{A}_t^P - \pi_t \mathcal{J}_t^P] + (1 - \hat{\pi}_t) [(1 - \pi_t) \mathcal{J}_t^O - \pi_t \mathcal{A}_t^O] = 0.$$

Weight the two projected first-order conditions accordingly and use (B.13). Terms involving ΔX_t cancel.

Expanding the remaining matrices gives

$$\begin{aligned} \mu_{r,t}^{BMG} &= e'_{BMG} [\gamma \Sigma_{r,t} + (1 - \gamma) b \Sigma_{rT,t}] X_t^M \\ &\quad + (1 - \gamma) b [\hat{\pi}_t \pi_t + (1 - \hat{\pi}_t) (1 - \pi_t)] e'_{BMG} (\Sigma'_{rT,t} - b \Sigma_{\Delta T}) X_t^M \\ &\quad + b e'_{BMG} [\hat{\pi}_t \pi_t \mu_{\Delta T}^P + (1 - \hat{\pi}_t) (1 - \pi_t) \mu_{\Delta T}^O] \\ &\quad - [\hat{\pi}_t \theta_P + (1 - \hat{\pi}_t) \theta_O] e'_{BMG} L_{r,t} \sigma_{\Delta T}. \end{aligned} \quad (\text{B.25})$$

Let $\mu_{r,t}^{BMG,0}$ be the right side evaluated at $\theta_P = \theta_O = 0$, with both subjective means set to $\mu_{\Delta T}$. Equation (B.8) now gives

$$\begin{aligned} \mathcal{B}_t^P &= \hat{\pi}_t [b \pi_t e'_{BMG} \sigma_{\Delta T} - e'_{BMG} L_{r,t} \sigma_{\Delta T}], \\ \mathcal{B}_t^O &= (1 - \hat{\pi}_t) [b (1 - \pi_t) e'_{BMG} \sigma_{\Delta T} - e'_{BMG} L_{r,t} \sigma_{\Delta T}], \\ \mu_{r,t}^{BMG} &= \mu_{r,t}^{BMG,0} + \mathcal{B}_t^P \theta_P + \mathcal{B}_t^O \theta_O. \end{aligned} \quad (\text{B.26})$$

Since $\theta_O \leq 0$, when $\mathcal{B}_t^O > 0$ this can be written as $\mathcal{B}_t^O [\bar{\theta}_t + \chi_t \theta_P - |\theta_O|]$, where $\bar{\theta}_t = \mu_{r,t}^{BMG,0} / \mathcal{B}_t^O$ and $\chi_t = \mathcal{B}_t^P / \mathcal{B}_t^O$. The BMG premium is therefore positive when $|\theta_O| < \bar{\theta}_t + \chi_t \theta_P$, negative when the inequality reverses, and zero at equality. This establishes all three cases of Property 1 on the stated parameter region. A sufficient condition for $\mathcal{B}_t^O > 0$ is $1 - \hat{\pi}_t > 0$ together with

$$b(1 - \pi_t) e'_{BMG} \sigma_{\Delta T} > e'_{BMG} L_{r,t} \sigma_{\Delta T}.$$

The condition states that the optimists' marginal damage loading exceeds their relative return-belief loading in this comparison. The threshold is economically attainable when it lies within the maintained range of optimists' tilts. ■

B.5 Proof of Property 2

Consider a nonempty family of interior equilibria of the allocation conditions, with long-only portfolios and fixed $\gamma > 1$, existing for all sufficiently large b . The family is maintained throughout this proof. Suppose that the approximation remains applicable and that, throughout this family,

$$e'_{BMG} \Sigma_{\Delta T} X_t^M \geq c > 0. \quad (\text{B.27})$$

Here c is a deterministic lower bound over the specified states and parameter values. Suppose also that constants $M_r, M_{rT}, M_\mu, M_\theta, M_{L\sigma}$ bound, respectively,

$$\|\Sigma_{r,t}\|, \quad \|\Sigma_{rT,t}\|, \quad \max_i \|\mu_{\Delta T}^i\|, \quad \max_i |\theta_i|, \quad \|L_{r,t} \sigma_{\Delta T}\|,$$

uniformly as b increases. Vector norms are Euclidean and matrix norms are the induced operator norms. Maintain these bounds uniformly over the states and parameter values considered.

Projecting (B.15)–(B.16), the leading damage-variance term is

$$(\gamma - 1)b^2[\pi_t^2 + (1 - \pi_t)^2]e'_{BMG} \Sigma_{\Delta T} X_t^M.$$

Full investment and long-only holdings imply $\|X_t^i\|, \|X_t^M\| \leq 1$ and $\|\pi_t^2 X_t^P + (1 - \pi_t)^2 X_t^O\| \leq 1$. Also $\pi_t^2 + (1 - \pi_t)^2 \geq 1/2$ and $\|e_{BMG}\| = \sqrt{2}$. Applying these bounds term by term yields

$$\mu_{r,t}^{BMG} \geq C_2 b^2 - C_1 b - C_0, \quad (\text{B.28})$$

where the following constants are introduced only to report the bound:

$$\begin{aligned} C_2 &= (\gamma - 1)c/2, \\ C_1 &= \sqrt{2}\{M_\mu + 2(\gamma - 1)M_{rT}\}, \\ C_0 &= \sqrt{2}\{\gamma M_r + M_\theta M_{L\sigma}\}. \end{aligned}$$

Since $C_2 > 0$, a sufficient threshold for the damage coefficient is

$$b > b_* \equiv \frac{C_1 + \sqrt{C_1^2 + 4C_2C_0}}{2C_2}. \quad (\text{B.29})$$

The positive quadratic damage-variance contribution then exceeds every adverse linear and bounded contribution. The same threshold applies to every state and belief configuration in the maintained family.

The bound c can be obtained directly from the covariance matrix and the admissible market weights. For $w = X_t^{M,B}$,

$$e'_{BMG} \Sigma_{\Delta T} X_t^M = w(\Sigma_{\Delta T}^{BB} - \Sigma_{\Delta T}^{BG}) + (1 - w)(\Sigma_{\Delta T}^{BG} - \Sigma_{\Delta T}^{GG}).$$

If w lies in a fixed interval, the minimum is attained at one of its endpoints. Both endpoint values must be positive. Under the common-factor specification $\Sigma_{\Delta T} = \sigma_{\Delta T} \sigma'_{\Delta T}$ and $\sigma_{\Delta T}^B > \sigma_{\Delta T}^G > 0$, one may take $c = (\sigma_{\Delta T}^B - \sigma_{\Delta T}^G) \sigma_{\Delta T}^G$. To translate (B.29) into a convexity threshold, use the damage coefficient

$$b(\eta) = \frac{\eta \delta \overline{\Delta T}^\eta}{1 + \delta \overline{\Delta T}^\eta}.$$

Assume the reference temperature is fixed at $\overline{\Delta T} \geq 1$ in the units used to specify the damage function. Then $b(\eta) \geq \eta \delta / (1 + \delta)$. Thus a sufficient, not necessarily sharp, threshold is

$$\bar{\eta} = \frac{1 + \delta}{\delta} b_*. \quad (\text{B.30})$$

For every $\eta > \bar{\eta}$, the lower bound is strictly positive. This proves Property 2 over the maintained feasible family. Economically, the marginal contribution of brown exposure to damage variance grows quadratically with b , while the bounded belief and covariance adjustments grow at most linearly. Required BMG compensation therefore becomes positive throughout the specified range of beliefs. ■

B.6 Lemma 2

Lemma 2 (Political Pressure and Endogenous Shares). Holding each type's climate expectation and $\overline{\Delta T}$ fixed, political pressure increases with the share of pessimists when both expected temperature changes are positive.

Proof. Differentiating the definition of ϕ_t gives

$$\frac{d\phi_t}{d\pi_t} = \frac{E_t^P[\Delta T_t] + E_t^O[\Delta T_t]}{\Delta \bar{T}} > 0. \quad (\text{B.31})$$

This proves the positive pressure response under the maintained fixed-expectation comparison. $\blacksquare$

B.7 Proof of Lemma 1

Timing and holding order. At date t , the realized portfolios, gross returns and common damage factor of generation $t - 1$ are observed before the share π_{t+1} is set. Assume $\beta > 0$, $U' > 0$, interior population shares, and $\Delta X_{t-1} > 0$ as specified by the demand-gap condition in B.1.

Suppose generation $t - 1$ enters with common wealth W_{t-1} , both types are fully invested, and the same realized damage factor $\bar{D}_{t-1}$ applies to their terminal wealth. Subtracting their wealths gives the exact identity

$$\hat{W}_t^P - \hat{W}_t^O = -W_{t-1} \bar{D}_{t-1} \Delta X_{t-1} (R_t^B - R_t^G). \quad (\text{B.32})$$

Here $\Delta X_{t-1} = X_{t-1}^{O,B} - X_{t-1}^{P,B}$ is signed. Full investment implies the green holding difference is its negative. The maintained positive holding gap gives $\Delta X_{t-1}^B = |\Delta X_{t-1}^G| = \Delta X_{t-1} > 0$. Write the logit rule in its equivalent form

$$\pi_{t+1} = \frac{\pi_t \exp(\beta \Delta U_{t-1})}{\pi_t \exp(\beta \Delta U_{t-1}) + 1 - \pi_t}, \quad \Delta U_{t-1} = U(\hat{W}_t^P) - U(\hat{W}_t^O).$$

At equal performance, its derivative with respect to ΔU_{t-1} is $\beta \pi_t (1 - \pi_t)$. Linearizing utility at the common reference wealth $\bar{W}_t$ and using (B.32) yields

$$\pi_{t+1} - \pi_t \simeq -\beta \pi_t (1 - \pi_t) U'(\bar{W}_t) W_{t-1} \bar{D}_{t-1} \Delta X_{t-1} (R_t^B - R_t^G). \quad (\text{B.33})$$

Under the joint first-order expansion around unit gross returns, $R_t^B - R_t^G \simeq r_t^{BMG}$. Therefore

$$\pi_{t+1} \simeq \pi_t - \Lambda_t r_t^{BMG}, \quad \Lambda_t = \beta \pi_t (1 - \pi_t) U'(\bar{W}_t) W_{t-1} \bar{D}_{t-1} \Delta X_{t-1}. \quad (\text{B.34})$$

All factors in Λ_t are positive under the maintained assumptions, so a higher realized BMG return lowers π_{t+1}

locally. The coefficient is measurable with respect to date- t information, and the entering cohort's share is fixed before date- $t + 1$ trading. Thus the timing of the adoption rule agrees with the valuation calculation. ■

B.8 Proof of Proposition 3 (Affine Valuation Equilibrium)

State and timing assumptions. Assume $s_t = (Emission_t^{cum}, \pi_t)'$ is a sufficient Markov state for conditional dividend growth, investment opportunities and political pressure. In particular, the realized performance input determining the next entering share is a function of s_t on the state space considered. Thus π_{t+1} is both measurable at date t and determined by the current state. Political-pressure inputs and inherited policy are also functions of s_t ; the cap is therefore a policy function of these two states.

Maintain the local transition

$$s_{t+1} = c + \Phi s_t + H u_{t+1}, \quad E_t[u_{t+1}] = 0,$$

with constant c , a 2×2 matrix Φ , and state-independent conditional second moments. The second row of $H u_{t+1}$ is zero almost surely under the adoption timing. For the affine calculation, fix the local coefficients $g_j, m_j, \sigma_j^2, \rho_j, c, \Phi$ and assume $0 < \rho_j < 1$ and $\rho_j \rho(\Phi) < 1$.

Coefficient matching. Taking expectations in the log-linear return identity and using the definition of $\mu_{r,t}^j$ gives

$$pd_t^j = k_j - r_f + E_t[\Delta d_{t+1}^j] + \rho_j E_t[pd_{t+1}^j] - \mu_{r,t}^j + \frac{1}{2} \sigma_j^2. \quad (\text{B.35})$$

Under finite conditional means and the transversality condition, forward iteration yields

$$pd_t^j = \sum_{s \geq 0} \rho_j^s E_t \left[k_j - r_f + \Delta d_{t+1+s}^j - \mu_{r,t+s}^j + \frac{1}{2} \sigma_j^2 \right].$$

Substitute $pd_t^j = a_j + B_j' s_t$, $E_t[\Delta d_{t+1}^j] = g_{0j} + g_j' s_t$, $\mu_{r,t}^j = m_{0j} + m_j' s_t$, and $E_t[s_{t+1}] = c + \Phi s_t$ into (B.35).

Matching state coefficients and constants gives

$$\begin{aligned} B_j'(I - \rho_j \Phi) &= g_j' - m_j', \\ a_j(1 - \rho_j) &= k_j - r_f + g_{0j} - m_{0j} + \rho_j B_j' c + \frac{1}{2} \sigma_j^2. \end{aligned} \quad (\text{B.36})$$

The spectral-radius condition implies convergence of $\sum_{s \geq 0} (\rho_j \Phi)^s$. Therefore

$$\begin{aligned} B'_j &= (g'_j - m'_j)(I - \rho_j \Phi)^{-1}, \\ a_j &= \frac{k_j - r_f + g_{0j} - m_{0j} + \rho_j B'_j c + \frac{1}{2} \sigma_j^2}{1 - \rho_j}. \end{aligned} \quad (\text{B.37})$$

These are the unique affine coefficients for the fixed local inputs, proving the coefficient-matching statement. The same restrictions imply the transversality limit for the affine solution.

For equilibrium existence, maintain the asset-level compatibility condition in B.1 and suppose the induced coefficient map is continuous and maps a nonempty compact convex set of admissible coefficients into itself. Brouwer's theorem then gives an equilibrium fixed point. When uniqueness of the joint coefficient system is required, additionally assume this map is a contraction on that set. Banach's theorem then gives uniqueness. The conditional linear solution and the joint fixed point have these respective existence and uniqueness conditions.

Loading signs. Write the entries of Φ as $\Phi_{EE}, \Phi_{E\pi}, \Phi_{\pi E}, \Phi_{\pi\pi}$. Direct inversion gives

$$\begin{aligned} B_{j,E} &= \frac{(g_{j,E} - m_{j,E})(1 - \rho_j \Phi_{\pi\pi}) + \rho_j \Phi_{\pi E}(g_{j,\pi} - m_{j,\pi})}{\det(I - \rho_j \Phi)}, \\ B_{j,\pi} &= \frac{\rho_j \Phi_{E\pi}(g_{j,E} - m_{j,E}) + (1 - \rho_j \Phi_{EE})(g_{j,\pi} - m_{j,\pi})}{\det(I - \rho_j \Phi)}. \end{aligned} \quad (\text{B.38})$$

Under the spectral-radius condition, both determinants are positive. For the signs used in the economic discussion, maintain a region in which the numerator of $B_{B,E}$ is negative, the numerator of $B_{B,\pi}$ is negative, the numerator of $B_{G,\pi}$ is positive, and the expression for $B_{G,E}$ exceeds that for $B_{B,E}$. These inequalities are restrictions on net dividend-growth and required-return responses, including their propagation through the two states.

An explicit sufficient benchmark is a common $\rho_B = \rho_G = \rho$, an uncoupled local transition $\Phi = \text{diag}(1, \Phi_{\pi\pi})$ with $|\rho \Phi_{\pi\pi}| < 1$, and

$$g_{B,E} - m_{B,E} < \min\{0, g_{G,E} - m_{G,E}\}, \quad g_{B,\pi} = g_{G,\pi} = 0, \quad m_{B,\pi} > 0 > m_{G,\pi}.$$

Then

$$B_{j,E} = \frac{g_{j,E} - m_{j,E}}{1 - \rho}, \quad B_{j,\pi} = -\frac{m_{j,\pi}}{1 - \rho \Phi_{\pi\pi}}. \quad (\text{B.39})$$

It follows that $B_{B,E} < B_{G,E}$, $B_{B,E} < 0$, and $B_{B,\pi} < 0 < B_{G,\pi}$. The factor $1/(1 - \rho)$ reflects unit persistence of cumulative emissions in this benchmark. The general two-state case uses (B.38); the uncoupled benchmark and the common discounted-loading benchmark in B.10 are separate comparisons, each with its stated assumptions. ■

B.9 Proof of Proposition 5

Local comparison. Maintain the timing, positive holding gap and local expansion in B.7, and the policy specification and positive slope in B.3. We isolate the effect of realized BMG performance transmitted through the entering cohort's belief share and the associated policy response. Other changes in state variables induced by the same realized return are held outside this channel. From Lemma 1, $d\pi_{t+1}/dr_t^{BMG} \simeq -\Lambda_t$ locally. Applying the chain rule and (B.22) at $t + 1$ gives

$$\left. \frac{d\mu_{r,t+1}^{BMG}}{dr_t^{BMG}} \right|_{\text{belief-policy}} = -\Lambda_t \left\{ \mathcal{S}_{t+1} \left[\Delta X_{t+1} - \frac{\psi \text{Cap}_{t+2}}{\text{Emission}_{t+2}^B - \text{Emission}_{t+2}^G} \frac{d\phi_{t+1}}{d\pi_{t+1}} \right] + \mathcal{R}_{t+1} \right\}. \quad (\text{B.40})$$

If $\Lambda_t > 0$ and $\mathcal{S}_{t+1} > 0$, this derivative is negative precisely when

$$\Delta X_{t+1} + \frac{\mathcal{R}_{t+1}}{\mathcal{S}_{t+1}} > \frac{\psi \text{Cap}_{t+2}}{\text{Emission}_{t+2}^B - \text{Emission}_{t+2}^G} \frac{d\phi_{t+1}}{d\pi_{t+1}}. \quad (\text{B.41})$$

This is reversal through the belief-policy channel. The reverse inequality gives momentum and equality gives no first-order response. With a zero holding gap the first-order adoption channel is inactive.

The required premium $\mu_{r,t+1}^{BMG}$ concerns returns from $t + 1$ to $t + 2$, while r_t^{BMG} is realized from $t - 1$ to t . The displayed derivative measures this specific transmission through the next entering cohort's composition, with other state responses held fixed. This is the belief-policy effect stated in Proposition 5. ■

B.10 Median Voting and the Comparison with Probabilistic Voting

Comparison assumptions. Maintain the reference-emissions policy specification in B.3. Fix the belief distribution and each type's climate expectations. Since the preferred cap is decreasing in expected warming, $\text{Cap}_{t+1}^P < \text{Cap}_{t+1}^O$. The median-voting rule selects the former for $\pi_t > 1/2$ and the latter for $\pi_t < 1/2$. Away

from the threshold its cap is constant with respect to the belief share. At the threshold a tie-breaking rule is required; derivatives across the switch are replaced by finite changes.

Write $\Delta Emission_{t+1} = Emission_{t+1}^B - Emission_{t+1}^G > 0$. Cap clearing gives

$$\Delta X_t^M = \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} e_{BMG}. \quad (B.42)$$

A higher cap increases the brown market weight and reduces the green market weight by the same amount. This comparison holds asset-level emissions fixed and requires feasible market weights under both rules.

Equation (B.25) gives BMG at either cap after eliminating individual portfolios. For the general comparison, evaluate $L_{r,t}$, its covariance matrices and $\hat{\pi}_t$ under each rule. The population-weighted mean-damage term in (B.15) is unchanged at fixed means, shares and temperature covariance, but the return-belief term and hedging adjustment need not be unchanged.

Now impose the benchmark $\rho B_E = \nu_t \mathbf{1}$, with the same ν_t under both rules. Hold b , α , Σ_K , $\Sigma_{\Delta T}$, the beliefs and π_t fixed across the comparison. Then

$$L_{r,t} = (1 - \alpha)I + \nu_t \mathbf{1} X_t^{M'}, \quad e'_{BMG} L_{r,t} = (1 - \alpha) e'_{BMG}, \quad L'_{r,t} X_t^M = (1 - \alpha + \nu_t) X_t^M.$$

It follows that

$$\begin{aligned} e'_{BMG} \Sigma_{r,t} e_{BMG} &= \alpha^2 e'_{BMG} \Sigma_K e_{BMG} + (1 - \alpha)^2 e'_{BMG} \Sigma_{\Delta T} e_{BMG}, \\ e'_{BMG} \Sigma_{rT,t} e_{BMG} &= (1 - \alpha) e'_{BMG} \Sigma_{\Delta T} e_{BMG}. \end{aligned}$$

Thus the four projected portfolio coefficients, $\mathcal{S}_t$ and $\hat{\pi}_t$ are independent of the cap in this benchmark. The mean-belief and return-belief terms in (B.25) are also unchanged. The remaining terms are linear in X_t^M , so their finite difference equals their slope times (B.42).

To identify that slope, note that $\mathcal{S}_t$ is the effective weighted average

$$\hat{\pi}_t (\mathcal{A}_t^P + \mathcal{J}_t^P) + (1 - \hat{\pi}_t) (\mathcal{A}_t^O + \mathcal{J}_t^O).$$

This follows either by inserting (B.24), or by differentiating the two portfolio equations at fixed coefficients.

Expanding the slope in (B.25) and using the preceding identities yields

$$\Delta\mu_{r,t}^{BMG} = \frac{\Delta Cap_{t+1}}{\Delta Emission_{t+1}} \left\{ \mathcal{S}_t + \nu_t [\gamma(1 - \alpha) + (1 - \gamma)b\{\hat{\pi}_t\pi_t + (1 - \hat{\pi}_t)(1 - \pi_t)\}] e'_{BMG} \Sigma_{\Delta T} e_{BMG} \right\}. \quad (\text{B.43})$$

This proves the finite cap-differential formula under common benchmark coefficients. The comparison maintains the same ν_t across regimes and feasible interior allocations throughout. If the expression in braces is positive, the BMG difference has the sign of the cap difference; if negative, the sign reverses; if zero, the BMG difference is zero. Equal caps also give zero difference. No sign for ν_t is needed for this algebra.

The cap ranking follows by substituting the political-pressure formula into the two voting rules. For positive expected warming, a pessimistic median gives $Cap_{t+1}^{MV} < Cap_{t+1}^{PV}$ for interior shares. For an optimistic median, median voting is stricter precisely when

$$\pi_t < \frac{2E_t^O[\Delta T_t]}{E_t^P[\Delta T_t] + E_t^O[\Delta T_t]},$$

with equality giving equal caps and the reverse inequality giving a more lenient median-voting cap. This condition is considered only within $\pi_t < 1/2$.

Finally, within a median-voting regime the counterpart of (B.22) is $d\mu_{r,t}^{BMG}/d\pi_t = \mathcal{S}_t \Delta X_t + \mathcal{R}_t$, with the derivatives in G_t^i evaluated holding that regime's cap fixed. Population shares still affect damage internalization. At an optimistic takeover, the pure policy jump uses $Cap^O - Cap^P > 0$ in (B.43) at fixed limiting shares and common coefficients. A finite change in the belief share also changes investor composition and incentives. Its total BMG effect combines the policy jump with those composition and incentive adjustments. ■